



**MVS/Extended Architecture
System-Data Administration**

Licensed
Program





MVS/Extended Architecture System-Data Administration

Licensed
Program

Third Edition (January 1987)

This edition replaces and makes obsolete the previous edition, GC26-4010-1.

This edition applies to Version 1 Release 1.2 of MVS/Extended Architecture Data Facility Product, Licensed Program 5665-284, and to any subsequent releases until otherwise indicated in new editions or technical newsletters.

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

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

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

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

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

© Copyright International Business Machines Corporation 1982, 1984, 1987

PREFACE

This manual, formerly titled MVS/Extended Architecture System Programming Library: Data Management, is now titled MVS/Extended Architecture System-Data Administration.

This publication provides information for system programmers about MVS/Extended Architecture Data Facility Product, and how to modify and extend the data management capabilities of the operating system.

ORGANIZATION

This publication contains the following chapters and appendixes:

- Chapter 1, "Managing the Volume Table of Contents (VTOC)" on page 1, defines and discusses the structure of the VTOC and VTOC index, and the use of system macros to read a data set control block, rename a data set, or delete a data set from the VTOC.
- Chapter 2, "Executing Your Own Channel Programs (EXCP)" on page 36, defines and discusses the use of the EXCP macro to control the organization of data based on device characteristics with your own channel programs.
- Chapter 3, "Reading from and Writing to Direct Access Devices (XDAP)" on page 75, defines and discusses the use of the XDAP macro to read, verify, and update blocks without using an access method.
- Chapter 4, "Password Protecting Data Sets" on page 84, defines and discusses system password protection and how to create and maintain the PASSWORD data set.
- Chapter 5, "Exit Routines" on page 96, defines and discusses some of the IBM-supplied exits for installation-written routines and authorized user programs.
- Chapter 6, "System Macro Instructions" on page 110, defines and discusses the system macros used to refer to, validate, and modify system data areas.
- Chapter 7, "Maintaining SYS1.IMAGELIB" on page 156, defines and discusses adding a UCS or FCB image to the system image library, and maintaining the UCS image tables.
- Chapter 8, "JES2 Support for the IBM 1403, 3203 Model 5, and 3211 Printers" on page 174, defines and discusses JES2 support for UCS alias names and the 3211 indexing feature.
- Chapter 9, "CATALOG, SCRATCH, and RENAME Dummy Modules" on page 176, defines and discusses the dummy modules for CATALOG, SCRATCH, and RENAME, and how to replace them.
- Chapter 10, "Specifying Buffer Numbers for DASD Data Sets" on page 177, defines and discusses the performance considerations when using the BUFNO keyword and subparameter.
- Appendix A, "CVAF - VTOC Access Macros" on page 178, defines and discusses the format of the VTOC access macros: CVAFDIR, CVAFDSM,
- Appendix B, "Examples of VTOC Access Macros" on page 196, defines and discusses examples of using the VTOC access macros in your programs.

- Appendix C, "VTOC Index Error Message and Associated Codes" on page 221, defines and discusses the error message and field codes issued by the Common VTOC Access Facility (CVAF).
- Appendix D, "Example of an OPEN Installation Exit Module" on page 225, defines and discusses a sample program listing for IFGOEX0B, the installation-written exit routine that takes control during OPEN for a DCB.

PREREQUISITE KNOWLEDGE

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

- Assembler language
- Standard program linkage conventions
- The utility programs IEHLIST and IEHPRGM
- Data management access methods and macro instructions

REQUIRED PUBLICATIONS

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

- Assembler H Version 2 Application Programming: Language Reference, GC26-4037, and Assembler H Version 2 Application Programming: Guide, GC26-4036, contain more information on coding in assembler language.
- MVS/Extended Architecture System Programming Library: Supervisor Services and Macro Instructions, GC28-1154, contains a description of standard linkage conventions.
- MVS/Extended Architecture Data Administration: Utilities, GC26-4018, describes how to use IEHLIST to maintain the VTOC, and IEHPRGM to protect data sets.
- MVS/Extended Architecture Data Administration Guide, GC26-4013, and MVS/Extended Architecture Data Administration: Macro Instruction Reference, GC26-4014, contain information on using access methods and macro instructions to do input and output.

Specific prerequisite reading is listed at the beginning of some chapters, as it relates to the particular topic.

RELATED PUBLICATIONS

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

Short Title (as it appears in the text)	Publication Title	Order Number
Access Method Services Reference	<u>MVS/Extended Architecture Integrated Catalog Administration: Access Method Services Reference</u>	GC26-4019
	<u>MVS/Extended Architecture VSAM Catalog Administration: Access Method Services Reference</u>	GC26-4075

Short Title (as it appears in the text)	Publication Title	Order Number
Assembler H V2 Application Programming: Guide	<u>Assembler H Version 2 Application Programming: Guide</u>	SC26-4036
Assembler H V2 Application Programming: Language Reference	<u>Assembler H Version 2 Application Programming: Language Reference</u>	GC26-4037
Catalog Administration Guide	<u>MVS/Extended Architecture Catalog Administration Guide</u>	GC26-4041
Checkpoint/Restart User's Guide	<u>MVS/Extended Architecture Checkpoint/Restart User's Guide</u>	GC26-4012
Conversion Notebook	<u>MVS/Extended Architecture Conversion Notebook</u>	GC28-1143
CVAF Diagnosis Reference	<u>MVS/Extended Architecture Common VTOC Access Facility Diagnosis Reference</u>	SY26-3929
DADSM and CVAF Diagnosis Guide	<u>MVS/Extended Architecture DADSM and Common VTOC Access Facility Diagnosis Guide</u>	SY26-3896
DADSM Diagnosis Reference	<u>MVS/Extended Architecture DADSM Diagnosis Reference</u>	SY26-3904
Data Administration Guide	<u>MVS/Extended Architecture Data Administration Guide</u>	GC26-4013
Data Administration: Macro Instruction Reference	<u>MVS/Extended Architecture Data Administration: Macro Instruction Reference</u>	GC26-4014
Debugging Handbook	<u>MVS/Extended Architecture Debugging Handbook, Volumes 1 through 5</u>	LC28-1164 ¹ LC28-1165 LC28-1166 LC28-1167 LC28-1168
Device Support Facilities User's Guide and Reference	<u>Device Support Facilities User's Guide and Reference</u>	GC35-0033
IBM System/370 XA Principles of Operation	<u>IBM System/370 Extended Architecture Principles of Operation</u>	SA22-7085
IBM 2821 Control Unit Component Description	<u>IBM 2821 Control Unit Component Description</u>	GA24-3312

Note:

¹ All five volumes may be ordered under one order number, LBOF-1015.

Short Title (as it appears in the text)	Publication Title	Order Number
IBM 3203 Printer Component Description and Operator's Guide	<u>IBM 3203 Printer Component Description and Operator's Guide</u>	GA33-1515
IBM 3211 Printer, 3216 Interchangeable Train Cartridge, and 3811 Printer Control Unit Component Description and Operator's Guide	<u>IBM 3211 Printer, 3216 Interchangeable Train Cartridge, and 3811 Printer Control Unit Component Description and Operator's Guide</u>	GA24-3543
IBM 3800 Printing Subsystem Programmer's Guide	<u>IBM 3800 Printing Subsystem Programmer's Guide</u>	GC26-3846
Initialization and Tuning	<u>MVS/Extended Architecture System Programming Library: Initialization and Tuning</u>	GC28-1149
JCL User's Guide	<u>MVS/Extended Architecture JCL User's Guide</u>	GC28-1351
JCL Reference	<u>MVS/Extended Architecture JCL Reference</u>	GC28-1352
Linkage Editor and Loader User's Guide	<u>MVS/Extended Architecture Linkage Editor and Loader User's Guide</u>	GC26-4011
Service Aids	<u>MVS/Extended Architecture System Programming Library: Service Aids</u>	GC28-1159
Supervisor Services and Macro Instructions	<u>MVS/Extended Architecture System Programming Library: Supervisor Services and Macro Instructions</u>	GC28-1154
System Logic Library	<u>MVS/Extended Architecture System Logic Library: Volume 8 of 17, Parts 1 and 2 (IOS)</u>	LY28-1234 (Part 1) LY28-1235 (Part 2)
System Macros and Facilities	<u>MVS/Extended Architecture System Programming Library: System Macros and Facilities, Volumes 1 and 2</u>	GC28-1150 GC28-1151
System Messages	<u>MVS/Extended Architecture Message Library: System Messages, Volumes 1 and 2</u>	GC28-1376 GC28-1377
System Modifications	<u>MVS/Extended Architecture System Programming Library: System Modifications</u>	GC28-1152

Short Title (as it appears in the text)	Publication Title	Order Number
TSO Command Language Reference	<u>MVS/Extended Architecture TSO Command Language Reference</u> (OS/VS2 TSO Command Language Reference, as updated by Supplement SD23-0259 for MVS/XA)	GC28-0646
TSO/E Command Language Reference	<u>MVS/Extended Architecture TSO Extensions TSO Command Language Reference</u>	SC28-1134
TSO/E Data Areas	<u>MVS/Extended Architecture TSO/E Data Areas</u> (plus Supplement LDB3-0276)	LYB8-1119
Utilities	<u>MVS/Extended Architecture Data Administration: Utilities</u>	GC26-4018
VSAM Administration: Macro Instruction Reference	<u>MVS/Extended Architecture VSAM Administration: Macro Instruction Reference</u>	GC26-4016

NOTATIONAL CONVENTIONS

A uniform system of notation describes the format of data management macro instructions. 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 these items may be selected.
- An ellipsis (...) indicates that multiple entries of the type immediately preceding the ellipsis are allowed.
- Other punctuation (parentheses, commas, spaces, and so forth) must be entered as shown. A space is indicated by a blank.
- **BOLDFACE** type indicates the exact characters to be entered, except as described in the bullets above. Such items must be entered exactly as illustrated.
- Lowercase underscored type specifies fields to be supplied by the user.
- **BOLDFACE UNDERSCORED** type indicates a default option. If the parameter is omitted, the underscored value is assumed.
- Parentheses () must enclose subfields if more than one is specified. If only one subfield is specified, you may omit the parentheses.

ADDRESS AND REGISTER CONVENTIONS

The following describes the meaning of each notation used to show how an operand can be coded:

symbol

The operand can be any valid assembler-language symbol.

(0)

General register 0 can be used as an operand. When used as an operand in a macro instruction, the register must be specified as the decimal digit 0 enclosed in parentheses as shown above.

(1)

General register 1 can be used as an operand. When used as an operand in a macro instruction, the register must be specified as the decimal digit 1 enclosed in parentheses as shown above. When you use register 1, the instruction that loads it is not included in the macro expansion.

(2-12)

The operand specified can be any of the general registers 2 through 12. All registers as operands must be coded in parentheses; for example, if register 3 is coded, it is coded as (3). When one of the registers 2 through 12 is used, it can be coded as a decimal digit, symbol (equated to a decimal digit), or an expression that results in a value of 2 through 12.

RX-Type Address

The operand can be specified as any valid assembler-language RX-type address. The following shows examples of each valid RX-type address:

Name	Operation	Operand
ALPHA1	L	1,39(4,10)
ALPHA2	L	REG1,39(4,TEN)
BETA1	L	2,ZETA(4)
BETA2	L	REG2,ZETA(REG4)
GAMMA1	L	2,ZETA
GAMMA2	L	REG2,ZETA
GAMMA3	L	2,=F'1000'
LAMBDA1	L	3,20(,5)

Both ALPHA instructions specify explicit addresses; REG1 and TEN have been defined as absolute symbols. Both BETA instructions specify implied addresses, and both use index registers. Indexing is omitted from the GAMMA instructions. GAMMA1 and GAMMA2 specify implied addresses. The second operand of GAMMA3 is a literal. LAMBDA1 specifies an explicit address with no indexing.

A-Type Address

The operand can be specified as any address that can be written as a valid assembler-language A-type address constant. An A-type address constant can be written as an absolute value, a relocatable symbol, or relocatable expression. Operands that require an A-type address are inserted into an A-type address constant during the macro expansion process. For more details about A-type address constants, see Assembler H Version 2 Application Programming: Language Reference.

absexp

The operand can be an absolute value or expression. An absolute expression can be an absolute term or an arithmetic combination of absolute terms. An absolute term can be a nonrelocatable symbol, a self-defining term, or the length attribute reference. For more details about absolute expressions, see Assembler H Version 2 Application Programming: Language Reference.

relexp

The operand can be a relocatable symbol or expression. A relocatable symbol or expression is one whose value changes by n if the program where it appears is relocated n bytes away from its originally assigned area of storage. For more details about relocatable symbols and expressions, see Assembler H Version 2 Application Programming: Language Reference.

SUMMARY OF CHANGES

RELEASE 1.2 LIBRARY UPDATE, JANUARY 1987

NEW PROGRAMMING SUPPORT

Support has been added for the conversion to ISO/ANSI/FIPS Version 3 tape labels, a new function of the WTOR installation exit.

Support has been added for 3480 block count checking, which compares the 3480 block count with the block count maintained by the system at end-of-volume.

Support has been added for the 3480 label processing PTF, which improves 3480 performance.

NEW DEVICE SUPPORT

Support has been added for the 3380 Mod AD4, BD4, AE4, and BE4.

Support has been added for the 3880 Model 21 and Model 23.

Support has been added for the 3480 Magnetic Tape Subsystem.

SERVICE CHANGES

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

RELEASE 1.2, FEBRUARY 1984

RESTRUCTURE AND UPDATING

This manual, formerly titled MVS/Extended Architecture System Programming Library: Data Management, is now titled MVS/Extended Architecture System-Data Administration.

Except for the updates noted below, the text of this manual is substantially the same as in Release 1.0. The following changes have been made:

- The Preface has been rewritten for stylistic consistency with the other MVS/XA documentation.
- The former Chapter 1, "Using Catalog Management Macro Instructions," has been moved to Catalog Administration Guide, and succeeding chapters have been renumbered.
- The "old" Appendix C, "Return Codes from VTOC Access Macros," has been combined with Appendix A, and the "Overviews" of the VTOC Access Macros formerly in Chapter 2 are now included in Appendix A.
- Chapter 7, "Maintaining SYS1.IMAGELIB," has been rewritten for ease of use.
- References to other manuals have been updated to reflect title changes for Release 1.2.

- Service changes have been made throughout the manual, and are indicated in the text by revision bars.

PROGRAMMING SUPPORT

Descriptions of the following have been added:

- ISO/ANSI/FIPS GDG password protection
- The new DADSM macro REALLOC, and REALLOC return codes



CONTENTS

Chapter 1. Managing the Volume Table of Contents (VTOC)	1
The VTOC	1
Data Set Control Block (DSCB) Format Types	2
Format-0 DSCB	3
Format-1 DSCB	3
Format-2 DSCB	3
Format-3 DSCB	3
Format-4 DSCB	4
Format-5 DSCB	4
Format-6 DSCB	4
Allocating and Releasing Space	4
The VTOC Index	6
An Example of a VTOC and Its Index	6
The VTOC Index Entry Record (VIER)	7
Contents of VIER Fields	7
Format of a VTOC Index Entry	8
When a VIER Is Created	9
A Tree of Linked VIERs	9
How to Find a Format-1 DSCB	9
Special Cases in a DSCB Search	9
The VTOC Pack Space Map (VPSM)	10
The VTOC Index Map (VIXM)	11
The VTOC Map of DSCBs (VMDS)	11
Structure of an Indexed VTOC	12
Scratch/Rename/Allocate Restrictions	12
Initializing and Maintaining the VTOC	13
Creating the VTOC and VTOC Index	13
Protecting the VTOC and VTOC Index	13
Resource Access Control Facility (RACF)	13
Authorized Program Facility (APF) Requirements	13
Password Protection	14
Copying/Restoring/Initializing the VTOC	14
Operations on Volumes Containing a Nonindexed VTOC	14
Operations on Volumes Containing an Indexed VTOC	14
Access the VTOC with DADSM Macros	15
Reading a Control Block from the VTOC	15
Reading a DSCB by Name (OBTAIN and CAMLST SEARCH)	15
Return Codes from OBTAIN	17
Reading a DSCB by Actual Device Address (OBTAIN and CAMLST SEEK)	17
Deleting a Data Set from the VTOC (SCRATCH and CAMLST SCRATCH)	18
Return Codes from SCRATCH	21
Renaming a Data Set in the VTOC (RENAME and CAMLST RENAME)	22
Return Codes from RENAME	24
Accessing the VTOC and its Index with CVAF Macros	25
Serialization and Updating	26
Identifying the Volume	26
Using Registers	26
Generating a CVPL (CVAF Parameter List)	27
Buffer Lists	28
Accessing the DSCB Directly	30
Accessing DSNs or DSCBs in Sequential Order	32
Obtaining Information from the VTOC Index	34
Diagnosing VTOC Errors	34
Actions Taken When an Error Occurs	34
Recovering from System or User Errors	34
GTF Trace	35
Listing a VTOC and VTOC Index	35
Chapter 2. Executing Your Own Channel Programs (EXCP)	36
Using EXCP in System and Problem Programs	37
How the System Uses EXCP	37
How To Use EXCP in Problem Programs	38
31-Bit IDAW Programming Notes	38
How EXCP Operates in a V=R Address Space	39
EXCP Requirements	39

Channel Program	39
Control Blocks	40
Input/Output Block (IOB)	40
Event Control Block (ECB)	40
Data Control Block (DCB)	40
Data Extent Block (DEB)	40
How the Channel Program Executes	41
Initiation of the Channel Program	41
Modification of a Channel Program during Execution	42
Completion of Execution	42
Interruption Handling and Error Recovery Procedures	42
Appendages	43
Start-I/O (SIO) Appendage	45
Program-Controlled Interruption (PCI) Appendage	45
End-of-Extent (EOE) Appendage	46
Channel-End (CHE) Appendage	47
Abnormal-End (ABE) Appendage	48
Making Your Appendages Part of the System	49
The Authorized Appendage List (IEAAPPO0)	49
Channel Programming Considerations	50
Macro Specifications for Use with EXCP	51
Defining Data Control Blocks for EXCP (DCB)	51
Foundation Block Parameters	52
EXCP Interface Parameters	52
Foundation Block Extension and Common Interface Parameters	54
Device-Dependent Parameters	55
DSORG Parameter of the DCBD Macro	58
Initializing Data Control Blocks (OPEN)	58
Executing a Channel Program (EXCP)	60
Assigning an Alternate Track and Copying Data from the Defective Track (ATLAS)	61
Using ATLAS	62
Operation of the ATLAS Program	63
Return Codes from the ATLAS Program	63
End of Volume (EOV)	65
Restoring Data Control Blocks (CLOSE)	66
Control Block Fields	67
Input/Output Block (IOB) Fields	67
Event Control Block (ECB) Fields	69
Data Extent Block (DEB) Fields	70
Executing Fixed Channel Programs in Real Storage (EXCPVR)	70
Building the List of Data Areas to be Fixed	71
Page Fix (PGFX) and Start-I/O (SIO) Appendage	72
Page-Fix List Processing	72
Chapter 3. Reading from and Writing to Direct Access Devices (XDAP)	75
XDAP Requirements	75
Macro Specifications for Use with XDAP	76
Defining a Data Control Block (DCB)	76
Initializing a Data Control Block (OPEN)	76
Executing Direct Access Programs (XDAP)	77
End of Volume (EOV)	79
Restoring a Data Control Block (CLOSE)	79
Control Blocks Used with XDAP	79
Event Control Block (ECB)	79
Input/Output Block (IOB)	79
Direct Access Channel Program	80
Converting a Relative Track Address to an Actual Track Address	80
Return Codes from the Conversion Routine	82
Converting an Actual Track Address to a Relative Track Address	82
Obtaining Sector Number of a Block on a Device with the RPS Feature	82
Chapter 4. Password Protecting Data Sets	84
Providing Data Set Security	84
PASSWORD Data Set Characteristics	86
Creating Protected Data Sets	86
Tape Volumes Containing More Than One Password-Protected Data Set	87
Protection Feature Operating Characteristics	87

Termination of Processing	87
Volume Switching	87
Data Set Concatenation	88
SCRATCH and RENAME Functions	88
Counter Maintenance	88
Maintaining the PASSWORD Data Set (PROTECT Macro)	88
PASSWORD Data Set Characteristics and Record Format (With PROTECT macro)	88
Number of Records for Each Protected Data Set	89
Protection-Mode Indicator	89
PROTECT Macro Syntax	90
PROTECT Macro Parameter Lists	91
Return Codes from the PROTECT Macro	95
Chapter 5. Exit Routines	96
DADSM Preprocessing and Postprocessing Exit Routines	96
The Exit Modules	96
The Exit Environment	96
When IGGPRE00 Gets Control	97
Rejecting a DADSM Request	97
Data that DADSM Passes to the Exits	98
Passing a Model Format-1 DSCB	99
When IGGPOST0 Gets Control	99
System Control Blocks	99
Registers at Entry to DADSM Exits	100
Registers at Return from DADSM exits	100
Return Codes from DADSM Exits	100
DCB OPEN Installation Exit	100
The Exit Module	101
The Exit Environment	101
Open Processing Before the DCB OPEN Exit Gets Control	101
Open Processing After the DCB OPEN Exit Gets Control	101
Getting Control from Open	102
Data that Open Passes to the Exit	102
Defaulting the DCB Buffer Number	103
Modifying the JFCB	103
Requesting Partial Release	104
Updating the Secondary Space Data	104
Registers at Entry to the DCB OPEN Exit	104
Registers at Return from the DCB OPEN Exit	105
Return Codes From the DCB OPEN Exit	105
Open/EOV Installation Exit for Format-1 DSCB Not Found	105
Data That OPEN/EOV Passes to the Exit	106
Registers at Entry to the Format-1 DSCB Not Found Exit	106
Registers at Return From the Format-1 DSCB Not Found Exit	106
Return Codes from the Format-1 DSCB Not Found Exit	107
Data Management Abend Installation Exit	107
Data that OPEN/EOV passes to the Exit	108
Registers at Entry to the Data Management ABEND Exit	109
Registers at Return from the Data Management ABEND Exit	109
Return Codes from the Data Management ABEND Exit	109
Chapter 6. System Macro Instructions	110
Introduction	110
Mapping System Data Areas	110
IEFUCBOB—Mapping the UCB	111
IEFJFCBN—Mapping the JFCB	111
CVT—Mapping the CVT	111
Obtaining I/O Device Characteristics	112
DEVTYPE Macro Specification	112
Device Characteristics Information	113
Return Codes from the DEVTYPE macro	115
Reading and Modifying a Job File Control Block	117
RDJFCB—Read a Job File Control Block	118
DEQ at Demount Facility for Tape Volumes	121
OPEN—Initialize Data Control Block for Processing the JFCB	122
High-Speed Positioning for the IBM 3480 Magnetic Tape Subsystem	123
Ensuring Data Security by Validating the Data Extent Block	124
DEBCHK—Macro Specification	125
Return Codes from the DEBCHK Macro	128
Purging and Restoring I/O Requests	128

PURGE—Halt or Finish I/O-Request Processing	130
Return Codes from PURGE Macro	132
Modifying the IOB Chain	132
RESTORE—Reprocess I/O Requests	133
Performing Track Calculations	133
TRKCALC—Standard Form	134
TRKCALC—Execute Form	137
TRKCALC—List Form	138
TRKCALC—DSECT Only	138
Input Register Usage for All Forms of 'MF'	138
Output from TRKCALC	138
Return Codes from TRKCALC	139
TRKCALC Macro Examples	139
Allocating a DASD Data Set	140
REALLOC—Execute Form	140
REALLOC—DSECT Only	142
REALLOC—List Form	142
Return Codes from REALLOC	142
Message Displays on the IBM 3480 Magnetic Tape Subsystem	143
MSGDISP—Displaying a Mount Message	144
MSGDISP—Displaying a Verify Message	146
MSGDISP—Displaying a Ready Message	147
MSGDISP—Displaying a Demount Message	148
MSGDISP—Resetting the Message Display	151
MSGDISP—Providing the Full Range of Display Options	152
Return Codes from MSGDISP	154
Chapter 7. Maintaining SYS1.IMAGELIB	156
UCS Images in SYS1.IMAGELIB	157
Adding a UCS Image to the Image Library	157
UCS Coding Examples	159
Verifying the UCS Image	162
Adding a UCS Image Name/Alias to a UCS Image Table	162
UCS Image Table Structure	162
Adding/Modifying a UCS Image Table Entry	166
Examples of Adding to the UCS Image Table	167
Example 1: Adding a New Band ID to the 4245 UCS Image Table (UCS5)	167
Example 2: Adding a New Default Entry to the 4248 UCS Image Table (UCS6)	168
Adding an FCB Image to the Image Library	169
Retrieving an FCB Image from SYS1.IMAGELIB	173
Chapter 8. JES2 Support for the IBM 1403, 3203 Model 5, and 3211 Printers	174
UCS Alias Names	174
The 3211 Indexing Feature	174
IBM 3203 Model 5 Printer	175
Chapter 9. CATALOG, SCRATCH, and RENAME Dummy Modules	176
Chapter 10. Specifying Buffer Numbers for DASD Data Sets Performance Considerations	177
Appendix A. CVAF - VTOC Access Macros	178
CVAFDIR Macro	178
Overview of the CVAFDIR Macro	178
Syntax	178
ACCESS: Read or Write a DSCB or VIR(S), or Release	
Buffer Lists	178
DSN: Specify the Name of the DSCB	179
BUFLIST: Specify One or More Buffer Lists	179
VERIFY: Verify that a DSCB is a Format-0 DSCB	179
UCB DEB: Specify the VTOC to Be Accessed	180
IOAREA: Keep or Free the I/O Work Area	180
MAPRCDS: Keep or Free MAPRCDS Buffer List and Buffers	181
IXRCDS: Retain VIERS in Virtual Storage	182
BRANCH: Specify the Entry to the Macro	183
MF: Specify the Form of the Macro	183
Return Codes from the CVAFDIR Macro	183
CVAFDSM Macro	185
Overview of the CVAFDSM Macro	185
Syntax	185

ACCESS=MAPDATA: Request Information from the Index Space Maps	185
MAP: Identify the Map to Be Accessed	185
EXTENTS: Identify Where Extents from the VPSM Are Returned	186
MAPRCDS: Keep or Free MAPRCDS Buffer List and Buffers	186
UCBIDEB: Specify the VTOC to Be Accessed	187
COUNT: Obtain a Count of Unallocated DSCBs or VIRs	187
CTAREA: Supply a Field to Contain the Number of Format-0 DSCBs	188
IOAREA: Keep or Free the I/O Work Area	188
BRANCH: Specify the Entry to the Macro	188
MF: Specify the Form of the Macro	189
Return Codes from the CVAFDSM Macro	189
CVAFSEQ Macro	190
Overview of the CVAFSEQ Macro	190
Syntax	190
ACCESS: Specify Relationship between Supplied and Returned DSN	190
BUFLIST: Specify One or More Buffer Lists	190
DSN: Specify Access by DSN Order or by Physical-Sequential Order	190
UCBIDEB: Specify the VTOC to Be Accessed	191
DSNONLY: Specify That Only the Data Set Name Be Read	191
ARG: Specify Where the Argument of the DSCB Is to Be Returned	191
IOAREA: Keep or Free the I/O Work Area	192
IXRCDS: Retain VIERS in Virtual Storage	192
BRANCH: Specify the Entry to the Macro	193
MF: Specify the Form of the Macro	193
Return Codes from the CVAFSEQ Macro	194
CVAFTST Macro	195
Overview of the CVAFTST Macro	195
Syntax	195
UCB: Specify the VTOC to Be Tested	195
Return Codes from the CVAFTST Macro	195
Appendix B. Examples of VTOC Access Macros	196
Example 1: Using the CVAFDIR Macro with an Indexed or Non-indexed VTOC	196
Example 2: Using the CVAFDIR Macro with an Indexed VTOC	201
Example 3: Using the CVAFSEQ Macro with an Indexed VTOC	206
Example 4: Using the CVAFSEQ Macro with a Nonindexed VTOC	211
Example 5: Using the CVAFTST and CVAFDSM Macros	217
Appendix C. VTOC Index Error Message and Associated Codes	221
Error Message	221
Explanation	221
System Action	221
Programmer Response	221
Routing and Descriptor Codes	221
Codes Put in the CVSTAT Field	222
Appendix D. Example of an OPEN Installation Exit Module	225
Processing in IFGOEXOB	225
Requesting Partial Release	225
Updating the Secondary Space Data	225
Index	235

FIGURES

1.	Locating the Volume Table of Contents (VTOC)	1
2.	Contents of VTOC—DSCBs Describing Data Sets	5
3.	Relationship of a VTOC to Its Index	6
4.	Format of the VTOC Index Entry Record (VIER)	8
5.	Structure of Linked VIERs	10
6.	An Index Map	11
7.	Format of a VTOC Map	12
8.	Format of the CVAF Parameter List (CVPL)	27
9.	CVFCTN Field of CVPL - Contents and Definitions	28
10.	Format of a Buffer List Header	29
11.	Format of a Buffer List Entry	30
12.	Entry Points, Returns, and Available Work Registers for Appendages	44
13.	Data Control Block (DCB) Format for EXCP (After OPEN)	53
14.	Input/Output Block (IOB) Format	68
15.	Event Control Block (ECB) after Posting of Completion Code (EXCP)	70
16.	Event Control Block (ECB) after Posting of Completion Code (XDAP)	80
17.	The XDAP Channel Programs	81
18.	Parameter List for Add Function	91
19.	Parameter List for REPLACE Function	92
20.	Parameter List for DELETE Function	93
21.	Parameter List for LIST Function	94
22.	Format of DADSM Preprocessing and Post-processing Exit Parameter List	98
23.	Format of DCB OPEN Installation Exit Parameter List (OIEXL)	102
24.	Output from DEVTYPE Macro	115
25.	Sample Code Using RDJFCB Macro	119
26.	Macro Definition, JCL, and Utility Statements for Adding PURGE Macro to the System Macro Library	129
27.	Macro Definition, JCL, and Utility Statements for Adding RESTORE Macro to the System Macro Library	129
28.	The PIRL and IOB Chain	133
29.	Sample Code to Add a 1403 UCS Image to SYS1.IMAGELIB	159
30.	Sample Code to Add a 3203 UCS Image to SYS1.IMAGELIB	160
31.	Sample Code to Add a 3211 UCS Image to SYS1.IMAGELIB	161
32.	UCS Image Table Entry Format	163
33.	UCS5 Image Table Contents	163
34.	UCS6 Image Table Contents	165
35.	Sample of the Standard FCB Image STD1	170
36.	Sample of the Standard FCB Image STD2	170
37.	Sample Code to Assemble and Add an FCB Load Module to SYS1.IMAGELIB	172

CHAPTER 1. MANAGING THE VOLUME TABLE OF CONTENTS (VTOC)

The direct access device storage management (DADSM) routines control allocation of space on direct access volumes through the volume table of contents (VTOC) of that volume, and through the VTOC index if one exists. This chapter gives an overview of the VTOC and the VTOC index and discusses how to use system macros to access the VTOC and VTOC index.

THE VTOC

The VTOC is a data set on a direct access volume that describes the contents of that volume. It resides in a single extent (that is, it is a continuous data set) anywhere on the volume after cylinder 0, track 0. Its address is located in the VOLVTOC field of the standard volume label (see Figure 1).

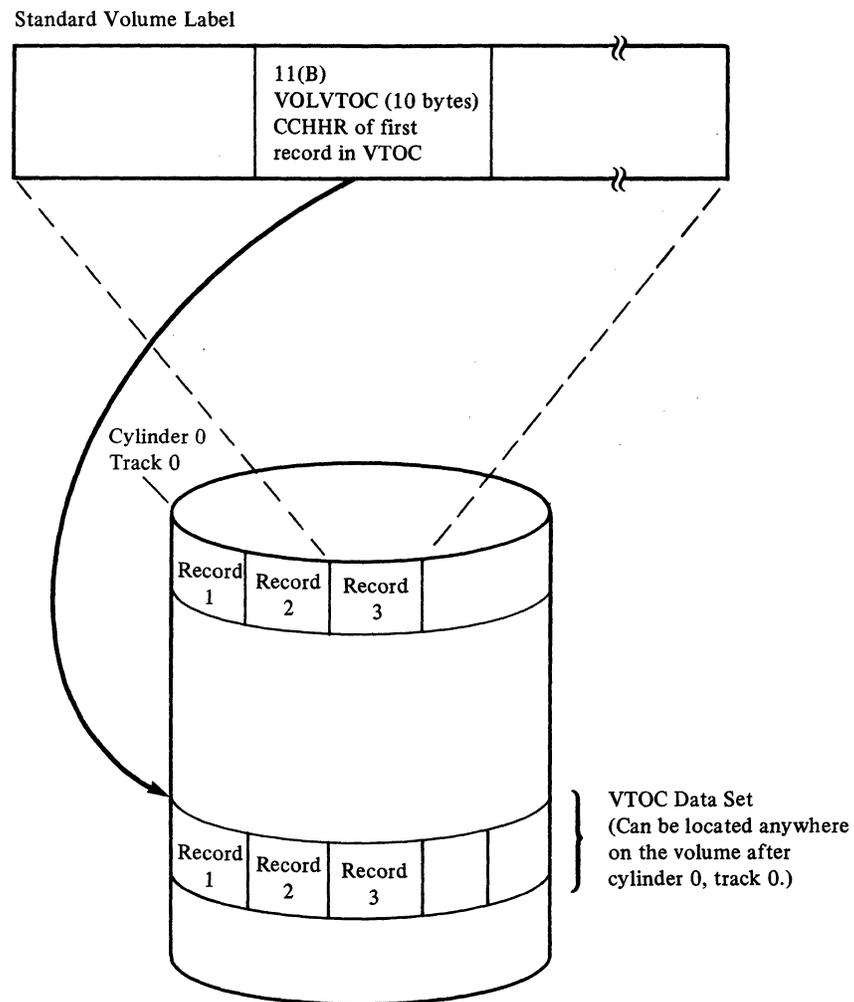


Figure 1. Locating the Volume Table of Contents (VTOC)

The VTOC is composed of 140-byte¹ data set control blocks (DSCBs) that correspond either to a data set or VSAM data space currently residing on the volume, or to contiguous, unassigned tracks on the volume. DSCBs for data sets or data spaces describe their characteristics. DSCBs for contiguous, unassigned tracks indicate their location.

DATA SET CONTROL BLOCK (DSCB) FORMAT TYPES

The VTOC has seven different kinds of DSCBs. This section lists the different kinds of DSCBs, what they are used for, how many exist on a volume, and how they are found.

The first record in every VTOC is the VTOC (format-4) DSCB that describes (1) the device that the volume resides on, (2) the attributes of the volume itself, and (3) the size and contents of the VTOC data set itself.

The format-4 DSCB is followed by a free-space (format-5) DSCB that, for a nonindexed VTOC, lists the extents on the volume that have not been allocated to a data set or VSAM data space. Each format-5 DSCB contains 26 extents. If there are more than 26 available extents on the volume, another format-5 DSCB will be built for every 26 extents. The format-5 DSCBs are chained, using the last field of each format-5 DSCB. An indexed VTOC does not use format-5 DSCBs for describing free space; however, one empty format-5 DSCB is provided to allow a basis for converting back to a nonindexed VTOC.

The third and subsequent DSCBs in the VTOC do not have any prescribed sequence.

A data set or VSAM data space is defined by one or more DSCBs in the VTOC of each volume on which it resides. The number of DSCBs needed to define a data set or VSAM data space is determined by (1) the organization of the data set (ISAM data sets need a format-2 DSCB to describe the index) and (2) the number of extents the data set or VSAM data space occupies (a format-3 DSCB is needed to describe the 4th through the 16th extents; additional format-3 DSCBs may be required to describe the extents for a VSAM data set cataloged in an integrated catalog facility catalog). Figure 2 on page 5 shows the general makeup of a VTOC and the DSCBs needed to define two types of data sets (ISAM and non-ISAM).

Data set A (in Figure 2 on page 5) is an ISAM data set; three DSCBs, a format-1, format-2, and format-3, are identified. Data sets B, C, and D could be sequential, partitioned, or direct data sets or they could be VSAM data spaces. Data set B has more than three extents and therefore requires both a format-1 and a format-3 DSCB.

Data sets C and D have three or fewer extents and need only a format-1 DSCB. The format-6 DSCB, pointed to by the format-4 DSCB, is used to keep track of the extents allocated in order to be shared by two or more data sets (split-cylinder data sets). For example, if data sets C and D share an extent made up of one or more cylinders, this extent would be described in the format-6 DSCB. Note that split-cylinder data sets cannot be allocated, but existing split-cylinder data sets can still be processed.

¹ The 140-bytes are defined as a 44-byte key portion followed by a 96-byte data portion. You may make references to the logical 140-byte DSCB or to either of its parts.

Format-0 DSCB

NAME: Free VTOC Record

FUNCTION: Describes an unused record in the VTOC (contains 140 bytes of binary zeros). To delete a DSCB from the VTOC, a format-0 DSCB is written over it.

HOW MANY: One for every unused 140-byte record on the VTOC. The DS4DSREC field of the format-4 DSCB is a count of the number of format-0 DSCBs on the VTOC. This field is not maintained for an indexed VTOC.

HOW FOUND: Search on key equal to X'00' (sometimes X'00000000') for a nonindexed VTOC; for an indexed VTOC, the VTOC map of DSCBs is used to find a format-0 DSCB.

Format-1 DSCB

NAME: Identifier

FUNCTION: Describes the first three extents of a data set or VSAM data space.

HOW MANY: One for every data set or data space on the volume, except the VTOC.

HOW FOUND: Search on key equal to the data set name. For an indexed VTOC, a CCHHR pointer for each data set name is in the VTOC index.

Format-2 DSCB

NAME: Index

FUNCTION: Describes the indexes of an ISAM data set.

HOW MANY: One for every ISAM data set (for a multivolume ISAM data set, a format-2 DSCB exists only on the first volume).

HOW FOUND: Chained from a format-1 DSCB that represents the data set.

Format-3 DSCB

NAME: Extension

FUNCTION: Describes the 4th through 16th extents of a data set or VSAM data space. Data sets and VSAM data spaces are restricted to 16 extents per volume. VSAM data sets cataloged in an ICF catalog may be extended to a maximum of 123 extents, in which case there may be as many as ten format-3 DSCBs.

HOW MANY: One for each data set or VSAM data space on the volume that has more than three extents. There may be as many as ten for a VSAM data set cataloged in an ICF catalog.

HOW FOUND: Chained from a format-2 or a format-1 DSCB that represents the data set or VSAM data space. In the case of a VSAM data set cataloged in an ICF catalog, the chain may be from a preceding format-3 DSCB.

Format-4 DSCB

NAME: VTOC

FUNCTION: Describes the extent and contents of the VTOC and provides volume and device characteristics. If the VTOC is indexed, certain fields of this DSCB are not maintained by DADSM. See "Structure of an Indexed VTOC."

HOW MANY: One on each volume.

HOW FOUND: VOLVTOC field of the standard volume label contains its address. It is always the first record in the VTOC.

Format-5 DSCB

NAME: Free Space

FUNCTION: On a nonindexed VTOC, describes the space on a volume that has not been allocated to a data set or to a VSAM data space (available space). For an indexed VTOC, format-5 is zero, and the volume pack space map describes the available space.

HOW MANY: One for every 26 noncontiguous extents of available space on the volume for a nonindexed VTOC; for an indexed VTOC, there is only one.

HOW FOUND: The first format-5 DSCB on the volume is always the second DSCB of the VTOC. If there is more than one format-5 DSCB, it will be chained from the previous format-5 DSCB via the DS5PTRDS field of each format-5 DSCB.

Format-6 DSCB

NAME: Shared Extent

FUNCTION: Describes the extents shared by two or more data sets (split-cylinder extents).

HOW MANY: One for every 26 split-cylinder extents on the VTOC.

HOW FOUND: The address of the first format-6 DSCB is contained in the DS4F6PTR field of the format-4 DSCB. If there is more than one format-6 DSCB on the volume, it will be chained from the previous format-6 DSCB via the DS6PTRDS field of the format-6 DSCB.

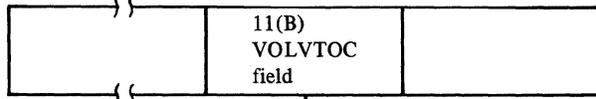
ALLOCATING AND RELEASING SPACE

The DADSM allocate and extend routines assign tracks and cylinders on direct access volumes for new data sets and VSAM data spaces. additional space for a data set or VSAM data space that has already exceeded its original, primary allocation. The DADSM scratch and partial release routines are used to release space that is no longer needed on a direct access volume.

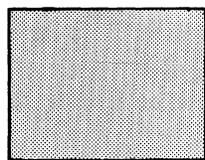
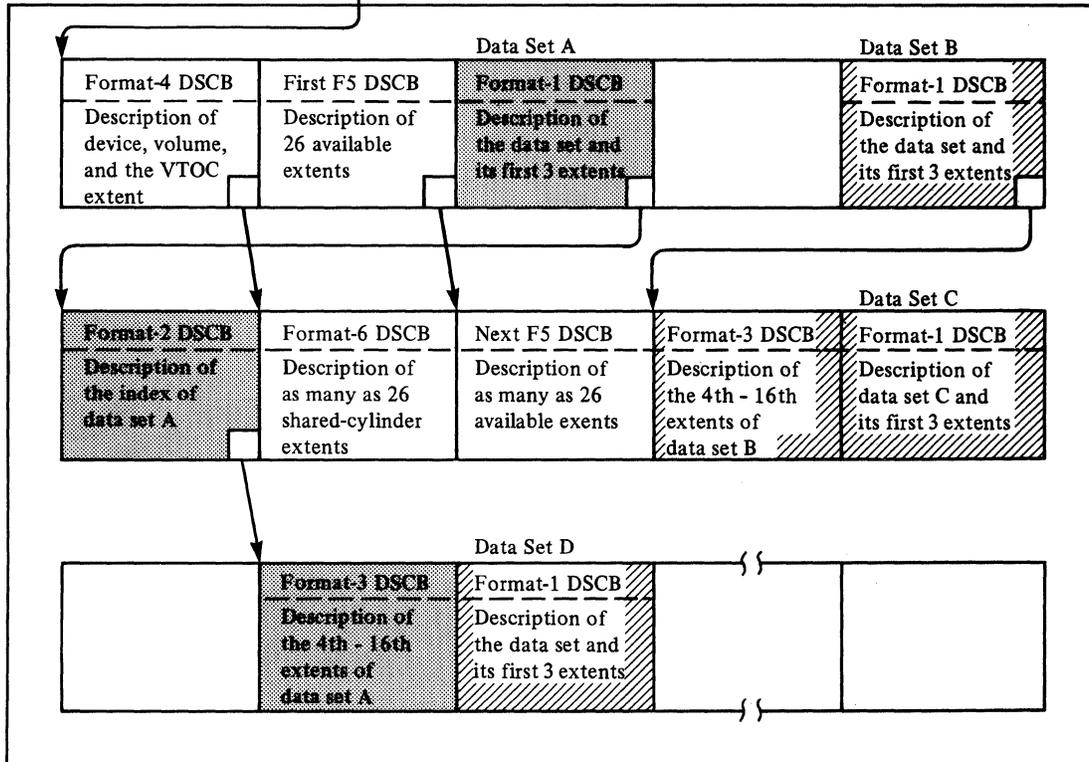
The DADSM routines allocate and release space by adding, deleting, and modifying the DSCBs. When space is needed on a volume, the allocate routines search the appropriate DSCBs for enough contiguous, available tracks to satisfy the request. If there are not enough contiguous tracks, the request is filled, using as many as five noncontiguous groups of free tracks. The appropriate DSCBs are modified to reflect the assignment of the tracks.

When space is released, the scratch routines free the DSCBs of the deleted data set or data space. For a nonindexed VTOC, to indicate that the tracks containing the affected data set or data space can be reallocated, a free space (format-5) DSCB is built (or modified if existent). For an indexed VTOC, the index is updated.

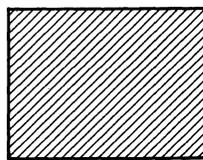
Standard Volume Label



VTOC Data Set



DSCB for an ISAM data set (Data Set A)



DSCB for a non-ISAM data set (Data Sets B, C, D) or a VSAM data space

Note: Empty boxes in the VTOC data set represent free VTOC Records (Format-0 DSCBs)

Figure 2. Contents of VTOC—DSCBs Describing Data Sets

THE VTOC INDEX

The VTOC index is a physical-sequential data set, residing on the same volume as the VTOC. It contains an index of data set names of format-1 DSCBs in the VTOC and free space information. The index is searched instead of the hardware keys.

The VTOC index is optional. You may build it when you initialize the volume, or for an existing VTOC (with the volume online or offline). You may subsequently inactivate it (online or offline) so that the VTOC is processed without using the index.

Each VTOC index is formatted by Device Support Facilities with physical blocks 2048 bytes in length. These physical blocks are the VTOC index records (VIRs), the basic structural units of the index. The kind of information they contain depends on the part of the index they belong to.

Several different kinds of records, each built from one or more VIRs, are in a VTOC index:

- The VTOC index entry record (VIER) that is used to access format-1 DSCBs and the format-4 DSCB
- The VTOC pack space map (VPSM) that shows what space has been allocated on a disk pack
- The VTOC index map (VIXM) that shows which VIRs have been allocated in the VTOC index
- The VTOC map of DSCBs (VMDS) that shows which DSCBs have been allocated in the VTOC

AN EXAMPLE OF A VTOC AND ITS INDEX

A format-1 DSCB in the VTOC contains the name and extent information of the VTOC index. The name of the index must be 'SYS1.VTOCIX.xxxxxxxx', where 'xxxxxxx' can be anything valid in a data set name and is generally the serial number of the volume containing the VTOC and its index. The name must be unique within the system to avoid ENQ contention. The relationship of a VTOC to its index is shown in Figure 3. Each of the components of the index is discussed separately in the following sections.

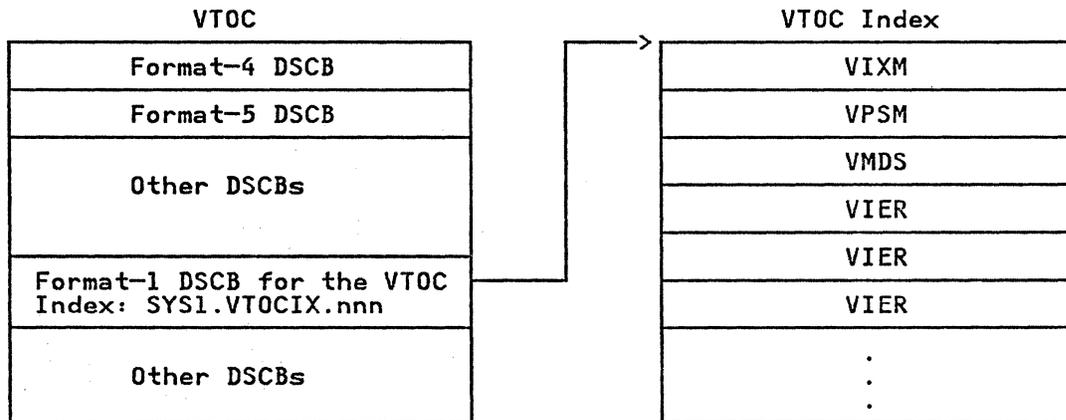


Figure 3. Relationship of a VTOC to Its Index

THE VTOC INDEX ENTRY RECORD (VIER)

VIERs have these characteristics:

- A VIER uses one VIR and contains variable-length index entries. The number of VIERs in an index vary depending upon the number of data sets on the volume.
- VIERs in a VTOC index may be on one or many levels. All index entries in a VIER are at the same index level. VIERs have a hierarchic relationship. Index entries in higher-level VIERs point to lower-level VIERs. Index entries in level-one VIERs (those at the lowest level) point to format-1 DSCBs for data sets on the volume.
- A higher-level VIER is created when the fourth lower-level VIER is created. When that new higher-level VIER is filled with pointers to lower-level VIERs, a new VIER at the same level is created. Again, when the fourth VIER at the same level is created, a VIER at a still higher level is created, adding another level to the index.

Contents of VIER Fields

Each VIER contains a header and sections (see Figure 4 on page 8). The VIER header contains:

- A field identifying the VTOC index record as a VIER.
- The relative byte address (RBA) of the VIER.
- A pointer to a VIER at the same level (hence, a "horizontal" pointer). The VIER pointed to contains index entries whose keys are greater than any key in the pointing VIER.
- The level number (LVL) of this VIER.
- The number (SECNO) of sections (a VIER contains eight sections).
- The length (SECL) of the sections (each section is 246 bytes in length).
- The offsets to the first-used and the last-used sections.
- The 44-byte high key of the VIER.

Each section contains:

- An offset to the last entry in the section (or zero if the section is empty)
- Index entries

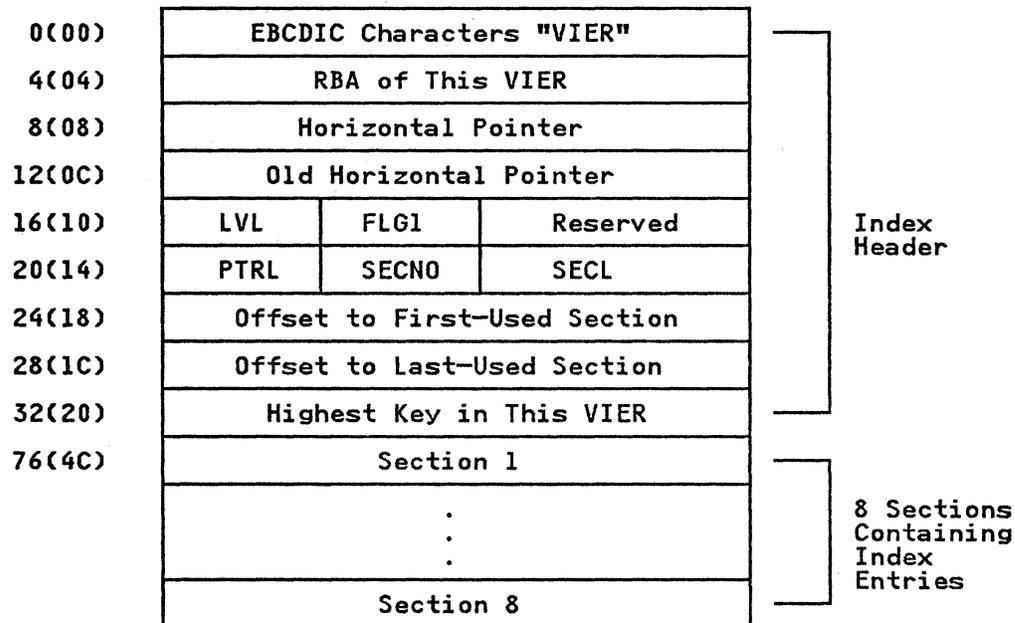


Figure 4. Format of the VTOC Index Entry Record (VIER)

Format of a VTOC Index Entry

The format of an index entry is:

FLG	KEYL	Unused	Record Pointer	Key
-----	------	--------	----------------	-----

Name	Offset	Bytes	Description
VXEFLG	00(00)	1	Flag byte
VXEKEYL	01(01)	1	Length of the VXEKEY field
VXEFC	02(02)	1	Unused
VXERPTR	03(03)	4 or 5	Record pointer
VXEKEY	07(07) or 08(08)	1 to 44	Name of a data set, if a level-one VIER; if not, the high key in the header of a lower-level VIER

Each index entry contains:

- A flag byte.
- A keylength field (containing a value of 1 to 44, depending on the length of the data set name).
- A record pointer (VXERPTR) that is one of the following:
 - In level-one VIERs, the 5-byte CCHHR of the format-1 or format-4 DSCB that represents the data set whose name is the key in the entry
 - In other VIERs, the 4-byte RBA of the lower-level VIER whose high key is the key in the entry

- A key that, for level 1 VIERS, is the data set name, and for level 2 or higher VIERS is the high key of a lower-level VIER. Trailing blanks are suppressed in the VTOC index entry.

When a VIER Is Created

The first level-one VIER is created when the VTOC index is created. Subsequent VIERS are created when a data set name is to be added to the VTOC index but the VIER where it should be added is full. A new VIER is created in the following manner:

- A new VIER is allocated.
- Half of the sections from a full VIER (those containing the highest keys) are moved into the new VIER, leaving each VIER half empty.
- The new index entry is added to one of the two VIERS, depending on its key.

A Tree of Linked VIERS

Figure 5 on page 10 shows how VIERS are related to each other. Note that the VIERS (which are simplified here—only the high key is shown in the header) form a type of "tree structure."

How to Find a Format-1 DSCB

In the search for the format-1 DSCB for a particular data set, one path along the tree structure is followed.

As seen in Figure 4 on page 8, a field in the header of a VIER contains the highest key of any index entry in that VIER. Beginning with this field in the first high-level VIER, the following search logic is used: Is the key of the data set (the data set name) lower than or equal to the VIER's high key? If neither, the test is again applied with the VIER having a greater high key pointed to by the horizontal pointer. This procedure continues until a VIER is found having a high key that is greater than or equal to the key of the data set. Comparisons are then made with the entries in the VIER's sections. Eventually, an entry is found with a key greater than or equal to the data set key. This entry points to a VIER at the next-lower level.

The search proceeds to successively lower levels until an entry in a level-two VIER is found whose key is greater than or equal to the key of the data set. This entry points to a level-one VIER that, in turn, contains an entry with a key that is equal to the data set key and that points to the format-1 DSCB for the desired data set.

Special Cases in a DSCB Search

If there is only one level in the VTOC index, the entries in the VIERS all point to format-1 DSCBs, so that only one level need be searched.

If an update to the VTOC index requires a new VIER and the update is interrupted (for example, because of an I/O error or a system failure), the entry in the level-n VIER may contain a key that is greater than the high key in the lower-level VIER pointed to by that entry. In this case, two VIERS at level n-1 may have to be searched. This situation is corrected when DADSM next processes the volume.

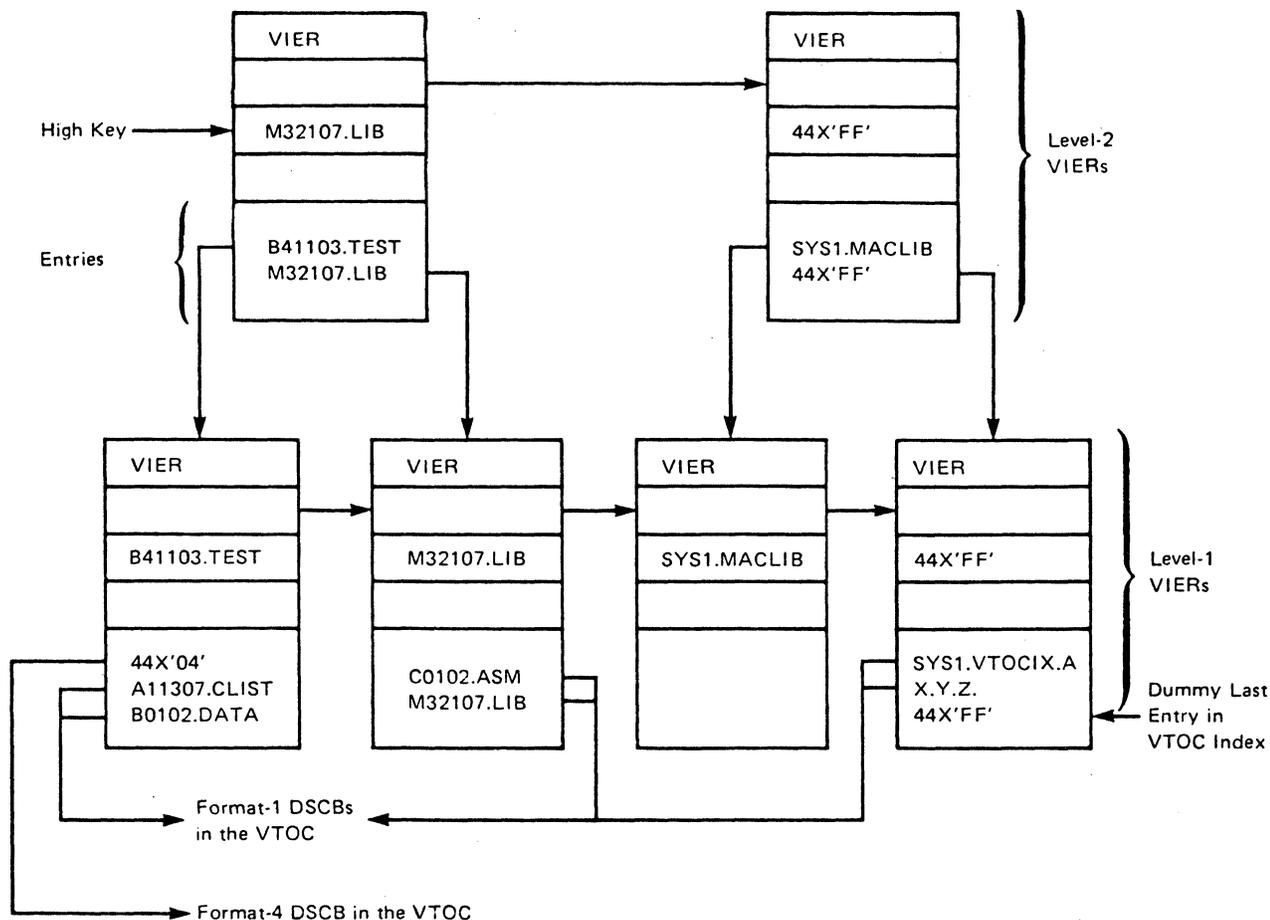


Figure 5. Structure of Linked VIERs

THE VTOC PACK SPACE MAP (VPSM)

The VPSM accounts for space on a disk pack. It shows what space on the volume has been allocated and what space remains free.

The map contains bit maps of the cylinders and tracks on the volume. A value of one indicates that the cylinder or track has been allocated; a value of zero, that it has not been allocated. The bit representing a cylinder is set to zero if no tracks on the cylinder have been allocated; it is set to one if any track has been allocated. Tracks assigned as alternate tracks are marked as allocated.

The VPSM replaces the chain of format-5 DSCBs, but one empty format-5 DSCB is left in the VTOC to allow for conversion back to a nonindexed VTOC, a process that requires reconstruction of a format-5 DSCB chain.

The format of an index map (including the VPSM) is shown in Figure 6 on page 11.

00(00)	ID of This Map		
04(04)	RBA of This Map		
08(08)	Horizontal Pointer to Next VIR		
12(0C)	Sequence Number of First Entry		
16(10)	VRFDA		VRFO
20(14)	FLG1	LUF1	LUOF
24(18)	Size of Large Unit Map		
28(1C)	SUF1	SUBIT	SUOF
32(20)	Size of Small Unit Map		
36(24)	Reserved		VIR
40(28)	RBA of First High-Level VIER		
	Large Unit Map (VTOC Pack Space Map Only)		
	Small Unit Map		
	VTOC Recording Facility Data (VTOC Index Map Only)		

Figure 6. An Index Map

THE VTOC INDEX MAP (VIXM)

The VIXM contains a bit map where each bit represents one VTOC index record (VIR). The status of the bit indicates whether the VIR is allocated (1) or unallocated (0).

An area of the VIXM is reserved for VTOC recording facility (VRF) data. (This is the facility that allows detection of and recovery from certain errors in an indexed VTOC.)

A field in the first VIXM record points to the first high-level VIER. Another field in the first VIXM record (VIR in Figure 7 on page 12) contains the number of VTOC index records that contain all the space maps.

THE VTOC MAP OF DSCBS (VMDS)

The VMDS contains a bit map where each bit represents one DSCB in the VTOC. The status of the bit indicates whether the DSCB is allocated (1) or unallocated (0).

Name	Offset	Bytes	Description
VIMAP	00(X'00')	2048	VTOC map
VIMH	00(X'00')	44	VTOC map header
VIMID	00(X'00')	4	Map ID in EBCDIC ('VPSM', 'VIXM', or 'VMDS')
VIMRBA	04(X'04')	4	RBA of this map
VIMHZPTR	08(X'08')	4	Horizontal RBA pointer to next VIR of this map
VIMORG	12(X'0C')	4	Sequence number of the first entry in the map
VIMVRFDA	16(X'10')	2	Offset to current VRF data (if VIMVRFSW=1) or offset where VRF data may be written (if VIMVRFSW=0), (first VIXM only)
VIMVRF0	18(X'12')	2	Offset to VRF area (first VIXM VIR only)
VIMFLG1	20(X'14')	1	Flag byte
VIMVRFSW	1... .. .xxx xxxxx		VRF data exists if 1 Reserved
VIMLUF1	21(X'15')	1	Large unit flag byte
VIMLUOF	22(X'16')	2	Offset into VIR of large unit map (zero if none)
VIMLUSZ	24(X'18')	4	Size in bits of large unit map
VIMSUF1	28(X'1C')	1	Small unit flag byte
VIMSUBIT	29(X'1D')	1	Number of small unit bits per large unit (zero if none)
VIMSUOF	30(X'1E')	2	Offset into VIR of small unit map
VIMSUSZ	32(X'20')	4	Size in bits of small unit map
	36(X'24')	3	Reserved
VIMVIR	39(X'27')	1	Number of map records (VIXM only)
VIMFHLV	40(X'28')	4	RBA of first high-level VIER (VIXM only)
VIMLUMAP	44(X'2C')	kk	Large unit map (kk is VIMLUSZ/8, rounded up)
VIMSUMAP	mm	nn	Small unit map (mm is VIMSUOF, nn is VIMSUSZ/8, rounded up)
VIMVRF	pp	qq	VRF area (pp is VIMVRF0, qq is remainder of first VIXM)

Figure 7. Format of a VTOC Map

STRUCTURE OF AN INDEXED VTOC

An indexed VTOC is identical to a nonindexed VTOC, except that, for an indexed VTOC, only a single format-5 DSCB exists and is empty, and certain format-4 DSCB data (the number of format-0 DSCBs and the CCHHR of the highest format-1 DSCB) is not maintained by DADSM. The DOS bit (bit 0 in field DS4VTOCI), set to one in the format-4 DSCB, indicates that these fields (and the format-5 DSCB) cannot be relied on. The index bit (bit 7 in field DS4VTOCI) is set in the format-4 DSCB; it indicates that a VTOC index exists.

SCRATCH/RENAME/ALLOCATE RESTRICTIONS

A VTOC index data set may not be scratched if the VTOC index is active. Neither may a VTOC index data set be renamed if the VTOC index is active, unless it is being renamed to another name beginning with 'SYS1.VTOCIX.'. A data set may not be renamed to a name beginning with 'SYS1.VTOCIX.' if there is already such a data set on the volume. Only one data set whose name begins with 'SYS1.VTOCIX.' may be allocated on a volume.

INITIALIZING AND MAINTAINING THE VTOC

CREATING THE VTOC AND VTOC INDEX

To prepare a volume for use (to initialize it), the Device Support Facilities utility is used. One of the things this utility does is to build the VTOC. After initialization, this VTOC will contain a format-4 DSCB and a format-5 DSCB. For a nonindexed VTOC, the format-5 DSCB contains an extent entry for all the free space on the volume; the initial number of extents in the format-5 DSCB is one or two, depending on where the VTOC is located on the volume. If the VTOC is located somewhere other than at the beginning or end of the volume, two extent entries are needed to describe the free space that precedes and follows it. For an indexed VTOC, the format-5 DSCB contains a zero.

A VTOC index can be created when a volume is initialized by using the Device Support Facilities command INIT and specifying the INDEX key word.

A nonindexed VTOC can be converted to an indexed VTOC by using the command BUILDIX and specifying the IXVTOC keyword. The reverse is also possible by using the BUILDIX command and specifying the OSVTOC keyword.

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

PROTECTING THE VTOC AND VTOC INDEX

Resource Access Control Facility (RACF)

You can protect the VTOC and VTOC index by using the Resource Access Control Facility (RACF). This is done by defining the volume serial entity under the RACF class DASDVOL. A user must be authorized to the DASDVOL/volume serial entity at the following levels:

- At the UPDATE level, to open the VTOC for output processing
- At the UPDATE level, to open for output processing any data set whose name begins with 'SYS1.VTOCIX.'
- At the ALTER level, to allocate, rename, or scratch any data set whose name begins with 'SYS1.VTOCIX.'
- At the ALTER level, to rename a data set to any name that begins with 'SYS1.VTOCIX.'

Neither the VTOC nor the VTOC index is protected from being opened for input processing by the DASDVOL/volume serial entity.

Note that neither the VTOC nor the VTOC index can be protected through the RACF class DATASET.

Authorized Program Facility (APF) Requirements

A program must be authorized by the authorized program facility (APF) to perform any of the following functions:

- Opening a VTOC for output processing
- Opening for output processing a data set whose name begins with 'SYS1.VTOCIX.'
- Allocating, renaming, or scratching any data set whose name begins with 'SYS1.VTOCIX.'

- Renaming a data set to any name that begins with 'SYS1.VTOCIX.'

Password Protection

The VTOC index data set may be password protected. The protection is the same as for any password-protected data set. Password checking is bypassed if the volume in which the VTOC index resides is protected by RACF through the DASDVOL class.

COPYING/RESTORING/INITIALIZING THE VTOC

OPERATIONS ON VOLUMES CONTAINING A NONINDEXED VTOC

- Restoring a Volume from a Dump Tape. There are no operational requirements if you change the volume serial number or do a partial restore that does not modify the VTOC. If you do a restore and change the VTOC size without changing the volume serial number, the volume must be varied offline after it is restored. You should not do a restore on a volume with an indexed VTOC.
- Copying a Volume. There are no operational requirements if you change the volume serial number or do not modify the VTOC of the receiving volume. If you do a copy and change the VTOC size without changing the volume serial number, the volume must be varied offline after it is copied. You should not do a copy from a volume with an indexed VTOC.

OPERATIONS ON VOLUMES CONTAINING AN INDEXED VTOC

You should use Device Support Facilities to convert a VTOC to a nonindexed format to update the volume. If you do not, take note of the following information:

- Initializing a Volume. If you do not change the volume serial number, the volume should be varied offline before starting the job.
- Restoring a Volume from a Dump Tape. There are no operational requirements if you change the volume serial number or do a partial restore that does not modify the VTOC or VTOC index. If you do a restore and modify the VTOC or VTOC index without changing the volume serial number, the volume should be varied offline after it is restored.
- Copying a Volume. There are no operational requirements if you change the volume serial number of the receiving volume or do a partial dump without modifying the VTOC or VTOC index. If you modify the VTOC or VTOC index without changing the volume serial number, the receiving volume should be varied offline after it is copied.
- Shared DASD Considerations. In shared DASD environments, whenever the VTOC index is modified or relocated or whenever the volume is changed from indexed VTOC to OS VTOC or from OS VTOC to indexed VTOC, the device should be varied offline to the sharing system or systems.

ACCESS THE VTOC WITH DADSM MACROS

You may use DADSM or CVAF to access the VTOC and its index. CVAF access is described in "Accessing the VTOC and its Index with CVAF Macros" on page 25. DADSM macros and associated tasks include:

- OBTAIN - Read a DSCB from a VTOC.
- REALLOC - DASD space allocation.
- RENAME - Rename a non-VSAM data set.
- SCRATCH - Release all space and DSCBs for a non-VSAM data set.

The REALLOC macro is described in "Allocating a DASD Data Set" on page 140.

This section tells how to use the OBTAIN, SCRATCH, and RENAME macro instructions. These macros are most commonly used by the operating system and the data set utility programs (IEHMOVE, IEBCOPY, and IEHPRGM), but you may use them in your own routines. The functions you can perform with these macros are:

- Reading a data set control block from the VTOC—OBTAIN
- Deleting a data set—SCRATCH
- Changing the name of a data set—RENAME

You can read a data set control block (DSCB) into virtual storage by using the OBTAIN and CAMLST macro instructions. There are two ways to specify the DSCB that you want to read: by using the name of the data set associated with the DSCB, or by using the absolute track address of the DSCB. You must provide a 140-byte data area in virtual storage, into which the DSCB will be read. When you specify the name of the data set, an identifier (format-1 or format-4) DSCB is read into virtual storage. To read a DSCB other than a format-1 or a format-4 DSCB, you must specify an absolute track address (see "Example" on page 18). (DSCB formats and field descriptions are contained in Debugging Handbook.)

You can delete a non-VSAM data set by using the SCRATCH and CAMLST macro instructions. This causes the DSCBs for the data set to be deleted.

You can change a data set name by using the RENAME and CAMLST macro instructions. This causes the data set name in the format-1 DSCB for the data set to be replaced with the new name.

Accompanying the descriptions of the macro instructions are coding examples, programming notes, and exception return code descriptions.

Note: OBTAIN, SCRATCH, and RENAME macro instructions cannot be used with a SYSIN or SYSOUT data set.

READING A CONTROL BLOCK FROM THE VTOC

Reading a DSCB by Name (OBTAIN and CAMLST SEARCH)

If you specify a data set name using OBTAIN and the CAMLST SEARCH option, the 96-byte data portion of the identifier (format-1) DSCB and the absolute track address of the DSCB are read into virtual storage. The absolute track address is a 5-byte field in the form CCHHR. The absolute track address field will contain zeros for VSAM and VIO data sets.

Because the VTOC does not contain a format-1 DSCB for a suballocated VSAM data space, an OBTAIN request, which searches the VTOC for such a data space's DSCB, fails. If the volume contains VSAM data sets, the OBTAIN routine uses information

from the VSAM catalog to build a pseudo format-1 DSCB, setting its CCHHR to zeros.

The format is:

<u>[symbol]</u> <u>listname</u>	OBTAIN CAMLST	<u>listname-addrx</u> SEARCH , <u>dsname-relexp</u> , <u>vol-relexp</u> , <u>wkarea-relexp</u>
------------------------------------	------------------	--

listname-addrx

points to the parameter list (labeled listname) set up by the CAMLST macro instruction.

SEARCH

this operand must be coded as shown.

dsname-relexp

specifies the virtual storage location of a fully qualified data set name. The area that contains the name must be 44 bytes long.

Note: A DSNAME of 44 bytes of X'04' (X'040404...04') can be used to read a format-4 DSCB.

vol-relexp

specifies the virtual storage location of the 6-byte volume serial number on which the DSCB is located.

wkarea-relexp

specifies the virtual storage location of a 140-byte work area that you must define.

Example: In the following example, the identifier (format-1) DSCB for data set A.B.C is read into virtual storage using the SEARCH option. The serial number of the volume containing the DSCB is 770655.

	OBTAIN	DSCBABC	READ DSCB FOR DATA
×			SET A.B.C INTO DATA
×			AREA NAMED WORKAREA
DSCBABC	CAMLST	SEARCH, DSABC, VOLNUM, WORKAREA	
DSABC	DC	CL44'A.B.C'	DATA SET NAME
VOLNUM	DC	CL6'770655'	VOLUME SERIAL NUMBER
WORKAREA	DS	140C	140-BYTE WORK AREA

Note: Check the return codes.

The OBTAIN macro instruction points to the CAMLST macro instruction. SEARCH, the first operand of CAMLST, specifies that a DSCB be read into virtual storage, using the data set name you have supplied at the address indicated in the second operand. DSABC, the second operand, specifies the virtual storage location of a 44-byte area into which you have placed the fully qualified name of the data set whose format-1 DSCB is to be read. VOLNUM, the third operand, specifies the virtual storage location of a 6-byte area into which you have placed the serial number of the volume containing the required DSCB. WORKAREA, the fourth operand, specifies the virtual storage location of a 140-byte work area into which the DSCB is to be returned.

Control will be returned to your program at the next executable instruction following the OBTAIN macro instruction. If the DSCB has been successfully read into your work area, register 15 will

contain zeros. Otherwise, register 15 will contain one of the following return codes. The return codes are shown in decimal, with hexadecimal values in parentheses.

Return Codes from OBTAIN

Code	Meaning
4(X'04')	The required volume was not mounted.
8(X'08')	The format-1 DSCB was not found in the VTOC of the specified volume.
12(X'0C')	A permanent I/O error was encountered, or an invalid format-1 DSCB was found when processing the specified volume, or an unexpected error return code was received from CVAF (Common VTOC Access Facility).
16(X'10')	Invalid work area pointer.

After execution of these macro instructions, the first 96 bytes of the work area contain the data portion of the identifier (format-1 or format-4) DSCB; the next 5 bytes contain the absolute track address (CCHHR) of the DSCB. These 5 bytes will contain zeros for VSAM or VIO data sets.

Reading a DSCB by Actual Device Address (OBTAIN and CAMLST SEEK)

You can read any DSCB from a VTOC using OBTAIN and the CAMLST SEEK option. You specify the SEEK option by coding SEEK as the first operand of the CAMLST macro and by providing the absolute device address of the DSCB you want to read, unless the DSCB is for a VIO data set. Only the SEARCH option can be used to read the DSCB of a VIO data set.

The format is:

<u>[symbol]</u> <u>listname</u>	OBTAIN CAMLST	<u>listname-addrx</u> SEEK , <u>cchhr-relexp</u> , <u>vol-relexp</u> , <u>wkarea-relexp</u>
------------------------------------	------------------	---

listname-addrx
points to the parameter list (labeled listname) set up by the CAMLST macro instruction.

SEEK
this operand must be coded as shown.

cchhr-relexp
specifies the virtual storage location of the 5-byte absolute device address (CCHHR) of a DSCB.

vol-relexp
specifies the virtual storage location of the 6-byte volume serial number on which the DSCB is located.

wkarea-relexp
specifies the virtual storage location of a 140-byte work area that you must define.

Example: In the following example, the DSCB at actual-device address X'00 00 00 01 07' is returned in the virtual storage location READAREA, using the SEEK option. The DSCB resides on the volume with the volume serial number 108745.

```

*           OBTAIN  ACTADDR           READ DSCB FROM
*           *                                           LOCATION SHOWN IN CCHHR
*           *                                           INTO STORAGE AT LOCATION
*           *                                           NAMED READAREA

ACTADDR    CAMLST  SEEK,CCHHR,VOLSER,READAREA
CCHHR      DC      XL5'0000000107' ABSOLUTE TRACK ADDRESS
VOLSER     DC      CL6'108745'   VOLUME SERIAL NUMBER
READAREA   DS      140C          140-BYTE WORK AREA

```

Note: Check the return codes.

The OBTAIN macro points to the CAMLST macro. SEEK, the first operand of CAMLST, specifies that a DSCB be read into virtual storage. CCHHR, the second operand, specifies the storage location that contains the 5-byte actual-device address of the DSCB. VOLSER, the third operand, specifies the storage location that contains the serial number of the volume where the DSCB resides. The fourth operand, READAREA, specifies the storage location to which the 140-byte DSCB is to be returned.

Control will be returned to your program at the next executable instruction following the OBTAIN macro instruction. If the DSCB has been successfully read into your work area, register 15 will contain zeros. Otherwise, register 15 will contain one of the following return codes. The return codes are shown in decimal, with hexadecimal values in parentheses.

Code	Meaning
4(X'04')	The required volume was not mounted.
8(X'08')	The format-1 DSCB was not found in the VTOC of the specified volume.
12(X'0C')	A permanent I/O error was encountered or an unexpected error return code was received from CVAF.
16(X'10')	Invalid work area pointer.
20(X'14')	The SEEK option was specified and the absolute track address (CCHHR) is not within the boundaries of the VTOC.

DELETING A DATA SET FROM THE VTOC (SCRATCH AND CAMLST SCRATCH)

You can use the SCRATCH and CAMLST macro instructions to delete a non-VSAM data set. SCRATCH processing deletes the associated data set control blocks (DSCBs) and makes the space occupied by the data set available for reallocation. Be aware that this process does not erase the data from the disk. Data sets that contain sensitive data should be erased (overwritten with zeros) before their space is made available.

If you want to scratch a data set being processed using virtual input/output (VIO), the data set must have been allocated for

use by your job. Scratching VIO data sets not allocated to your job is not allowed.

If the data set to be deleted is sharing one or more cylinders with one or more data sets (a split-cylinder data set), the space will not be made available for reallocation until all data sets on the shared cylinders are deleted.

A data set cannot be deleted if the expiration date in the identifier (format-1) DSCB has not passed, unless you choose to ignore the expiration date. You specify that the expiration date is to be ignored by using the OVRD option in the CAMLST macro instruction.

For information on RACF-defined data sets, see RACF General Information Manual. You may scratch a RACF-defined data set (that is, the DSCB indicates RACF-defined) only if you have alter access authority to either the data set/volume serial in the DATASET class, or to the volume serial in the DASDVOL class (if the volume is RACF-defined).

If a data set to be deleted is stored on more than one volume, either a device must be available for mounting the volumes or at least one volume must be mounted. In addition, all other required volumes must be serially mountable.

When deleting a data set, you must build a volume list in virtual storage. This volume list consists of an entry for each volume on which the data set resides. The first two bytes of the list indicate the number of entries in the list. Each 12-byte entry consists of a 4-byte device code, a 6-byte volume serial number, and a 2-byte scratch status code that should be initialized to zero. Device codes are presented in Debugging Handbook in the description of UCBTYP.

If the space to be deleted is a VSAM data space, you must use the DELETE command provided by access method services. For complete information about the DELETE command, see Access Method Services Reference.

Volumes are processed in the order that they appear in the volume list. The volume at the beginning of the list is processed first. If a volume is not mounted, a message is issued to the operator requesting that a volume be mounted. (A volume mount message will not be issued for a mass storage system (MSS) virtual volume; however, a status code will be returned to your program.) This is only done if register 0 has been loaded with the UCB associated with the device where unmounted volumes are to be mounted. (The device must be allocated to your job.) If you do not load register 0 with a UCB address, its contents must be zero, and at least one of the volumes in the volume list must be mounted before the SCRATCH macro instruction is issued.

If the requested volume cannot be mounted, the operator issues a reply indicating that the request cannot be fulfilled. A status code is then set in the last byte of the volume pointer (the second byte of the scratch status code) for the unavailable volume, and the next volume indicated in the volume list is processed.

The format is:

<u>[symbol]</u> <u>listname</u>	SCRATCH CAMLST	<u>listname-addrx</u> SCRATCH , <u>dsname-relexp</u> ,, <u>vol list-relexp</u> [, <u>OVRD</u>]
------------------------------------	-------------------	---

listname-addrx
points to the parameter list (labeled listname) set up by the CAMLST macro instruction.

SCRATCH

this operand must be coded as shown.

dsname-relexp

specifies the virtual storage location of a fully qualified data set name. The area that contains the name must be 44 bytes long. The name must be defined by a C-type define constant (DC) instruction.

vol list-relexp

specifies the virtual storage location of an area that contains a volume list. The area must begin on a halfword boundary.

OVRD

when coded as shown, specifies that the expiration date in the DSCB should be ignored.

Example: In the following example, data set A.B.C is deleted from two volumes. The expiration date in the identifier (format-1) DSCB is ignored.

	SR	0,0	SET REG 0 TO ZERO
	SCRATCH	DELABC	DELETE DATA SET A.B.C
*			FROM TWO VOLUMES,
*			IGNORING EXPIRATION
*			DATE IN THE DSCB
DELABC	CAMLST	SCRATCH,DSABC,,VOLIST,,OVRD	
DSABC	DC	CL44'A.B.C'	DATA SET NAME
VOLIST	DC	H'2'	NUMBER OF VOLUMES
	DC	X'3030200E'	3380 DISK DEVICE CODE
	DC	CL6'000017'	VOLUME SERIAL NO.
	DC	H'0'	SCRATCH STATUS CODE
	DC	X'3030200E'	3380 DISK DEVICE CODE
	DC	CL6'000018'	VOLUME SERIAL NO.
	DC	H'0'	SCRATCH STATUS CODE

Note: Check the return codes and SCRATCH status codes.

The SCRATCH macro instruction points to the CAMLST macro instruction. SCRATCH, the first operand of CAMLST, specifies that a data set be deleted. DSABC, the second operand, specifies the virtual storage location of a 44-byte area where you have placed the fully qualified name of the data set to be deleted. VOLIST, the fourth operand, specifies the virtual storage location of the volume list you have built. OVRD, the sixth operand, specifies that the expiration date in the DSCB of the data set to be deleted be ignored.

When you attempt to delete a password-protected data set that is not also RACF-protected, the operating system issues a message (IEC301A) to ask the operator at the console or the terminal operator of a remote console to enter the password. The data set will be scratched only if the password supplied is associated with a WRITE protection mode indicator. The protection mode indicator is described in Chapter 5, "Password Protecting Data Sets."

Control is returned to your program at the next executable instruction following the SCRATCH macro instruction. If the data set has been successfully deleted, register 15 will contain zeros, and the scratch status code in the volume list entry for each volume will be set to zero. Otherwise, register 15 will contain one of the return codes that follow. To determine whether the data set has been successfully deleted from each volume on which it resides, you must examine the scratch status code, that is, the last byte of each entry in the volume list.

Return Codes from SCRATCH

Code	Meaning
4(X'04')	No volumes containing any part of the data set were mounted, nor did register 0 contain the address of a unit that was available for mounting a volume of the data set. The data set may be a VIO data set that was not allocated during your job. (This return code is accompanied by a scratch status code of 5 in each entry of the volume list.)
8(X'08')	An unusual condition was encountered on one or more volumes.
12(X'0C')	The volume list passed was invalid. The scratch status code (the last byte of each volume list entry) will not have been modified during scratch processing.

After the SCRATCH macro instruction is executed, the last byte of each 12-byte entry in the volume list indicates one of the following conditions in binary codes:

Scratch
Status
Code

Meaning

0	All DSCBs for the data set have been deleted from the VTOC on the volume pointed to.
1	The VTOC of this volume does not contain the format-1 DSCB for the data set to be deleted.
2	The macro instruction failed when the correct password was not supplied in the two attempts allowed, or an attempt was made to scratch a VSAM data space or data set cataloged in an ICF catalog.
3	The data set was not deleted from this volume because either the OVRD option was not specified or the retention cycle has not expired.
4	A permanent I/O error was encountered, or an invalid format-1 DSCB was found when processing this volume, or an unexpected error return code was received from CVAf.
5	It could not be verified that this volume was mounted, and no device was available for mounting this volume.
6	The operator was unable to mount this volume. For Mass Storage Systems (MSS), a volume mount failure occurred. For a JES3-managed virtual volume, JES3 would not allow the volume to be mounted.
7	The specified data set could not be scratched because it was being used.
8	The DSCB indicates the data set is defined to RACF, but either the user is not authorized to access the data set or the volume, or the data set is a VSAM data space, or the data set is cataloged in an ICF catalog, or the data set is not defined to RACF.

| RENAMING A DATA SET IN THE VTOC (RENAME AND CAMLST RENAME)

You rename a data set that is not cataloged in an ICF or VSAM catalog by using the RENAME and CAMLST macro instructions. These cause the data set name in all format-1 DSCBs for the data set to be replaced by the new name you supply. (VIO data sets cannot be renamed.)

If a data set to be renamed is stored on more than one volume, either a device must be available for mounting the volumes, or at least one volume must be mounted. In addition, all other volumes of the data set must be serially mountable.

For information on RACF-defined data sets, see RACF General Information Manual. Only a user with alter access authority may rename a RACF-defined data set.

When renaming a data set, you must build a volume list in virtual storage. This volume list consists of an entry for each volume on which the data set resides. The first two bytes of the list indicate the number of entries in the list. Each 12-byte volume list entry consists of a 4-byte device code, a 6-byte volume serial number, and a 2-byte rename status code that should be initialized to zero. Device codes are presented in Debugging Handbook. Volumes are processed in the order that

they appear in the volume list. The first volume on the list is processed first. If a volume is not mounted, a message is issued to the operator requesting that the volume be mounted. (A volume mount message will not be issued for an MSS volume; however, a status code will be returned to your program.) This is only done if you indicate the direct access device on which unmounted volumes are to be mounted by loading register 0 with the address of the UCB associated with the device to be used. (The device must be allocated to your job.) If you do not load register 0 with a UCB address, its contents must be zero, and at least one of the volumes in the volume list must be mounted before the RENAME macro instruction is executed.

If the operator cannot mount a volume in the volume list, a reply is issued that the request cannot be fulfilled. A status code is then set in the last byte of the volume list entry (the second byte of the rename status code) for the unavailable volume, and the next volume indicated in the volume list is processed or requested.

The format is:

[symbol] <u>listname</u>	RENAME CAMLST	<u>listname-addrx</u> RENAME , <u>dsname-relexp</u> , <u>new name-relexp</u> , <u>vol list-relexp</u>
-----------------------------	------------------	---

listname-addrx

points to the parameter list (labeled listname) set up by the CAMLST macro instruction.

RENAME

this operand must be coded as shown.

dsname-relexp

specifies the virtual storage location of a fully qualified data set name to be replaced. The area that contains the name must be 44 bytes long. The name must be defined by a C-type define constant (DC) instruction.

new name-relexp

specifies the virtual storage location of a fully qualified data set name that is to be used as the new name. The area that contains the name must be 44 bytes long. The name must be defined by a C-type Define Constant (DC) instruction.

vol list-relexp

specifies the virtual storage location of an area that contains a volume list. The area must begin on a halfword boundary.

Example: In the following example, data set A.B.C is renamed D.E.F. The data set resides on two volumes.

	SR	0,0		SET REG 0 TO ZERO
	RENAME	DSABC		CHANGE DATA SET
				NAME A.B.C TO D.E.F
DSABC	CAMLST	RENAME,OLDNAME,NEWNAME,VOLIST		
OLDNAME	DC	CL44'A.B.C'		OLD DATA SET NAME
NEWNAME	DC	CL44'D.E.F'		NEW DATA SET NAME
VOLIST	DC	H'2'		TWO VOLUMES
	DC	X'3030200E'		3380 DISK DEVICE CODE
	DC	CL6'000017'		VOLUME SERIAL NO.
	DC	H'0'		RENAME STATUS CODE
	DC	X'3030200E'		3380 DISK DEVICE CODE
	DC	CL6'000018'		VOLUME SERIAL NO.
	DC	H'0'		RENAME STATUS CODE

Note: Check the return codes and RENAME status codes.

The RENAME macro instruction points to the CAMLST macro instruction. RENAME, the first operand of CAMLST, specifies that a data set be renamed. OLDNAME, the second operand, specifies the virtual storage location of a 44-byte area where you have placed the fully qualified name of the data set to be renamed. NEWNAME, the third operand, specifies the virtual storage location of a 44-byte area where you have placed the new name of the data set. VOLIST, the fourth operand, specifies the virtual storage location of the volume list you have built.

Control is returned to your program at the next executable instruction following the RENAME macro instruction. If the data set has been successfully renamed, register 15 will contain zeros, and the rename status code in the volume list entry for each volume will be set to zero. Otherwise, register 15 will contain one of the return codes below. To determine whether the data set has been successfully renamed on each volume where it resides, you must examine the rename status code, the last byte of each entry in the volume list.

Return Codes from RENAME

Code	Meaning
4(X'04')	No volumes containing any part of the data set were mounted, nor did register 0 contain the address of a unit that was available for mounting a volume of the data set to be renamed. The data set may be a VIO data set and cannot be renamed. (This return code is accompanied by a rename status code of 5 in each entry of the volume list.)
8(X'08')	An unusual condition was encountered on one or more volumes.
12(X'0C')	The volume list passed was invalid. The rename status code, the last byte of each volume list entry, will not have been modified during rename processing.

After the RENAME macro instruction is executed, the last byte of each 12-byte entry in the volume list indicates one of the following conditions in binary code:

Rename Status Code	Meaning
0	The format-1 DSCB for the data set has been renamed in the VTOC on the volume pointed to.
1	The VTOC of this volume does not contain the format-1 DSCB for the data set to be renamed.
2	The macro instruction failed when the correct password was not supplied in the two attempts allowed, or the user tried to rename a VSAM data space or VSAM data set cataloged in an ICF catalog.
3	A data set with the new name already exists on this volume.
4	A permanent I/O error was encountered, or an invalid format-1 DSCB was found when trying to rename the data set on this volume, or an unexpected error return code was received from CVAF.
5	It could not be verified that the volume was mounted, and no device was available for mounting the volume.
6	The operator was unable to mount this volume. For Mass Storage Systems (MSS), a volume mount failure occurred. For a JES3-managed virtual volume, JES3 would not allow the volume to be mounted.
7	The specified data set could not be renamed on this volume because it was being used.
8	The data set is defined to RACF but either the user is not authorized to alter the data set or the data set is defined to RACF on multiple volumes.

When you attempt to rename a password-protected data set, the operating system issues a message (IEC301A) to ask the operator or remote console operator to verify the password. The data set will be renamed only if the password supplied is associated with a WRITE protection mode indicator. The protection mode indicator is described in Chapter 4, "Password Protecting Data Sets" on page 84.

ACCESSING THE VTOC AND ITS INDEX WITH CVAF MACROS

You may use CVAF or DADSM to access the VTOC or its index. DADSM access is described in "Access the VTOC with DADSM Macros" on page 15.

CVAF macros and associated tasks include:

- CVAFDIR—Directly access DSCBs or VTOC index records.
- CVAFDSM—Obtain volume free space information.
- CVAFSEQ—Retrieval of the following:
 - Data set names from an active VTOC index.
 - DSCBs in physical-sequential order.
 - DSCBs in data set name order (index required).
- CVAFTST—Determine if a DASD volume has an active VTOC index.

Appendix A, "CVAF - VTOC Access Macros" on page 178 contains detailed descriptions of these macros. Appendix B, "Examples of VTOC Access Macros" on page 196 contains examples of their use.

Serialization and Updating

CVAF requires that you provide all necessary system resource serialization for your request. You can ensure the integrity of multiple data elements (sets of DSCBs and/or VIRs) returned by CVAF only when you provide adequate resource serialization. This exposure is compounded if multiple requests are required for the compilation of a desired set of DSCBs and/or VIRs. You must weigh possible system performance loss due to serialization against the potential loss of data integrity.

Updates made without adequate serialization may compromise the integrity of the volume's VTOC, the VTOC index, and/or any associated data set.

CVAF honors requests to modify the volume's VTOC and/or index for authorized programs only. CVAF assumes that an authorized program holds an exclusive RESERVE (or ENQ) on the qname (major name) of SYSVTOC, rname (minor name) of the volume's serial number, with the scope of SYSTEMS.

The SYSVTOC qname does not serialize access to the format-1 DSCB for a data set. You may provide this serialization by allocating the data set with disposition OLD, MOD, or NEW (not SHR). This causes the proper ENQ, ensuring that no other job can update that data set's format-1 DSCB.

Identifying the Volume

If you are authorized, you may identify the volume to the CVAFDIR, CVAFDSM, and CVAFSEQ macros by specifying the address of its UCB. If you are not authorized, you must identify the volume by specifying the address of a SAM or EXCP DEB opened to the volume's VTOC.

The DEB can be obtained by opening a DCB using the RDJFCB and OPEN TYPE=J macros. The DCB's DDNAME is that of a DD statement allocated to the unit whose VTOC is to be accessed. After issuing the RDJFCB macro, the JFCBDSNM field is overlaid with the data set name of the format-4 DSCB: 44X'04'. The DCB is opened for INPUT using OPEN TYPE=J. The DEB address is in DCB field, DCBDEBA. The OPEN macro is described under the section "OPEN—Initialize Data Control Block for Processing the JFCB" on page 122; the RDJFCB macro is described under "RDJFCB—Read a Job File Control Block" on page 118.

If a CVAF macro call has specified IOAREA=KEEP, a subsequent CVAF call using a different CVPL may omit the UCB and DEB keywords and supply the IOAREA address from the other CVPL. You can use the IOAREA keyword to do this.

The above does not apply to the CVAFTST macro. Only a UCB may be supplied to identify the VTOC, and no authorization is required.

Using Registers

Register 1 contains the address of the CVAF parameter list (CVPL). Register 15 contains the return code when processing for a function is complete.

Generating a CVPL (CVAF Parameter List)

The CVAFTST macro expands to provide its only parameter (UCB address) in register 1, and calls the associated CVAF module. All of the remaining CVAF macros use the CVPL to pass parameters to CVAF. CVAF uses the CVPL to return information related to the CVAF request.

CVAF generates a CVPL when you specify the CVAFDIR, CVAFDSM, or CVAFSEQ macro with MF=L or MF=I as a subparameter. If you do not specify the MF subparameter, MF=I is the default. The CVIIVT bit indicates whether an indexed or nonindexed VTOC is being accessed. The CVSTAT field contains feedback when an error occurs. The address of the map records buffer list is returned in the CVMRCDS field. The address of the VIER buffer list is returned in the CVIRCDS field. The CVAF I/O area address is returned in the CVIOAR field.

You may use the CVPL generated by the MF=L or MF=I form of the CVAFDIR, CVAFDSM, or CVAFSEQ macro (by using the MF=E keyword) to execute a different function from the one specified by the macro that originally generated the CVPL.

The ICVAFPL macro maps the CVPL. The format of the CVPL is shown in Figure 8.

Name	Offset	Bytes	Description
CVPL			
CVLBL	00(X'00')	4	EBCDIC "CVPL"
CVLTH	04(X'04')	2	
CVFCTN	06(X'06')	1	Function Byte (See Figure 9 on page 28)
CVSTAT	07(X'07')	1	Status Information
CVFL1	08(X'08')	1	First Flag Byte
CVIIVT		1... ..	Indexed VTOC Accessed
CV1IOAR		.1..	IOAREA=KEEP
CV1PGM		..1.	BRANCH=(YES,PGM)
CV1MRCDS		...1	MAPRCDS=YES
CV1IRCDS	 1...	IXRCDS=KEEP
CV1MAPIX	1..	MAP=INDEX
CV1MAPVT	1.	MAP=VTOC
CV1MAPVL	1	MAP=VOLUME
CVFL2	09(X'09')	1	Second Flag Byte
CV2HIVIE		1...	HIVIER=YES
CV2VRF		.1..	VRF data exists
CV2CNT		..1.	COUNT=YES
CV2RCVR		...1	RECOVER=YES
CV2SRCH	 1...	SEARCH=YES
CV2DSNLY	1..	DSNONLY=YES
CV2VER	1.	VERIFY=YES
CV2NLEVL	1	New highest level VIER (output)
CVUCB	12(X'0C')	4	UCB address
CVDSN	16(X'10')	4	Data set name address
CVBUFL	20(X'14')	4	Buffer list address
CVIRCDS	24(X'18')	4	Index VIRs buffer list address
CVMRCDS	28(X'1C')	4	Map VIRs buffer list address
CVIOAR	32(X'20')	4	I/O area address
CVDEB	36(X'24')	4	DEB address
CVARG	40(X'28')	4	Argument address
CVSPACE	44(X'2C')	4	SPACE parameter list address
CVEXTS	48(X'30')	4	Extent table address
CVBUFL2	52(X'34')	4	New VRF VIXM buffer list address
CVVRFDA	56(X'38')	4	VRF data address
CVCTAR	60(X'3C')	4	Count area address

Figure 8. Format of the CVAF Parameter List (CVPL)

The possible contents of the CVFCTN field in the CVPL and their meanings are as follows:

Name	Description	
CVDIRD	X'01'-CVAFDIR	ACCESS=READ
CVDIWR	X'02'-CVAFDIR	ACCESS=WRITE
CVDIRLS	X'03'-CVAFDIR	ACCESS=RLSE
CVSEQGT	X'04'-CVAFSEQ	ACCESS=GT
CVSEQGTE	X'05'-CVAFSEQ	ACCESS=GTEQ
CVDMIXA	X'06'-CVAFDSM	ACCESS=IXADD
CVDMIXD	X'07'-CVAFDSM	ACCESS=IXDLT
CVDMALC	X'08'-CVAFDSM	ACCESS=ALLOC
CVDMRLS	X'09'-CVAFDSM	ACCESS=RLSE
CVDMMAP	X'0A'-CVAFDSM	ACCESS=MAPDATA
CVVOL	X'0B'-CVAFVOL	ACCESS=VIBBLD
CVVRFRD	X'0C'-CVAFVRF	ACCESS=READ
CVVRFWR	X'0D'-CVAFVRF	ACCESS=WRITE

Figure 9. CVFCTN Field of CVPL - Contents and Definitions

Buffer Lists

A buffer list consists of one or more chained control blocks, each with a header and buffer list entries. The header indicates whether the buffer list is for DSCBs or VTOC index records. The entries point to and describe the buffers.

Buffer lists can be created in two ways:

- Directly, when you fill in the arguments and buffer addresses of DSCBs or VIRs to be read or written
- Indirectly (by CVAF), when you code the IXRCDS=KEEP and/or MAPRCDS=YES keywords

CVAF buffer lists are mapped by the ICVAFBFL macro. Figure 10 on page 29 shows the format of a buffer list header. Figure 11 on page 30 shows the format of a buffer list entry.

BUFFER LIST HEADER: The buffer list header indicates whether the buffer list describes buffers for DSCBs or for VTOC index records. The DSCB bit must be set to one and the VIR bit to zero in order for CVAF to process a request to read or write a DSCB. CVAF requires that you provide buffer lists and buffers in the caller's protect key. CVAF uses the protect key and subpool fields in the buffer list header only if you code ACCESS=RLSE.

Each buffer list header contains a count of the number of entries in the buffer list that directly follows the header.

The forward chain address is used to chain buffer lists together. DSCB buffer lists must not be chained to VIR buffer lists, and VIR buffer lists must not be chained to DSCB buffer lists.

The format of the buffer list header is shown in Figure 10.

Name	Offset	Bytes	Description
BFLHDR	0(X'00')	8	Buffer list header
BFLHNOE	0(X'00')	1	Number of entries
BFLHFL	1(X'01')	1	Flag byte and key
BFLHKEY		xxxx	Protect key of buffer list and buffers
BFLHVIR	 1...	Buffer list entries describe VIRs
BFLHDSCB	1..	Buffer list entries describe DSCBs
	xx	Reserved
	2(X'02')	1	Reserved
BFLHSP	3(X'03')	1	Identifies the subpool of buffer list and buffers
BFLHFCHN	4(X'04')	4	Forward chain address of next buffer list

Figure 10. Format of a Buffer List Header

BUFFER LIST ENTRY: A buffer list contains one or more entries. Each entry provides the buffer address, the length of the DSCB or VIR buffer, the argument, and an indication whether the argument is an RBA, a TTR, or a CCHHR.

The fields and bit uses are listed below.

- For a VIR buffer, the TTR and CCHHR bits must be 0, and the RBA bit must be 1.
- For a DSCB buffer, the RBA bit must be 0, and one of either the TTR or CCHHR bits must be set to 1 (they must not both be 1).
- The BFLEAUPD bit is an output indicator from CVAF that the BFLEARG field of a VIR buffer list was updated.
- The BFLEMOD bit indicates that a VIR buffer was modified and must be written; if no BFLEMOD bits are on in any of the entries for a CVAFDIR ACCESS=WRITE, all buffers are written.
- The BFLESKIP bit is used to cause an entry to be ignored.
- The BFLEIOER bit is an output indicator from CVAF to indicate an I/O error occurred during reading or writing of the DSCB or VIR.
- The BFLELTH field is the length of the buffer; for a DSCB buffer, the length must be 96 or 140; for a VIR buffer, the length must be the length of the buffer divided by 256.
- The BFLEARG field is the argument of the DSCB or VIR. Specify the desired format of the 5-byte field by setting either the BFLECHR, BFLETTR, or BFLERBA bit to 1. The respective BFLEARG values and formats are as follows:
 - CCHHR=5 byte CCHHR
 - TTR=0TTR0
 - RBA=One byte of 0 followed by a 4-byte RBA

The optional and required values for BFLEARG are dependent upon the variables associated with a given request. These are described in the following topics.

The format of the buffer list entry is shown in Figure 11.

Name	Offset	Bytes	Description
BFLE	0(X'00')	12	Buffer list entry
BFLEFL	0(X'00')	1	Flag byte
BFLEBBA		1... ..	Argument is RBA
BFLECHR		.1.. ..	Argument is CCHHR
BFLETRR		..1.	Argument is TTR
BFLEAUPD		...1	CVAF updated argument field
BFLEMOD	 1...	Data in buffer has been modified
BFLESKIP	1..	Skip this entry
BFLEIOER	1.	I/O error
	x	Reserved
	1(X'01')	1	Reserved
BFLELTH	2(X'02')	1	Length of VIR buffer divided by 256 or length of DSCB buffer
BFLEARG	3(X'03')	5	Argument of VIR or DSCB
BFLEATTR	4(X'04')	3	TTR of DSCB
BFLEARBA	4(X'04')	4	RBA of VIR
BFLEBUF	8(X'08')	4	Buffer address

Figure 11. Format of a Buffer List Entry

Accessing the DSCB Directly

CVAFDIR may be used to read or write a DSCB. CVAFDIR may also be used to read or write VTOC index records for indexed VTOCs. "CVAFDIR Macro" on page 178 discusses detailed information about the CVAFDIR VTOC access macro.

After a CVAFDIR call, the CVAF parameter list bit, CVLIVT, may be tested to determine whether the VTOC is indexed or nonindexed.

SPECIFYING A DATA SET NAME TO READ OR WRITE A DSCB: If you want to read or write a single DSCB by specifying only the data set name (that is, BFLEARG is zero), you must specify either ACCESS=READ or ACCESS=WRITE.

Specify the address of the data set name in the DSN keyword, then specify the address of the buffer list in the BUFLIST keyword. Each of these areas and the associated buffers must be in your protect key.

The buffer list must contain at least one buffer list entry with the skip bit off and a pointer to a 96-byte buffer. You must not provide 140-byte buffers. You may chain buffer lists together, but CVAF uses only the first eligible entry.

For an indexed VTOC, the index will be searched for the data set name and, if it is found, the DSCB argument obtained will be put in the buffer list entry and used to read or write the DSCB. If the data set name is not found in the index, a key search of the VTOC will be performed.

For a nonindexed VTOC, a channel program will be used to do a key search of the VTOC to locate the data set name and read or write the DSCBs. If the data set name is found, the DSCB argument will be put in the buffer list entry.

The DSCB argument returned in the buffer list entry will be in the format determined by the buffer list entry bits BFLECHR or BFLETRR.

If the data set name is not found in the VTOC, register 15 will contain a return code of 4, and CVSTAT will contain an error code of 1.

SPECIFYING THE DSCB LOCATION: If you want to read or write a DSCB by specifying the DSCB's location (that is, BFLEARG), you must specify either ACCESS=READ or ACCESS=WRITE.

Specify the address of the data set name in the DSN keyword, then specify the address of the buffer list in the BUFLIST keyword. Each of these areas and the associated buffer(s) must be in your program's protect key.

The buffer list must have at least one buffer list entry with the skip bit off and a pointer to a 96-byte or 140-byte buffer. You may chain buffer lists together, but CVAF uses only the first eligible entry.

If the buffer is for a 96-byte read or write, CVAF issues a channel program to verify that the key in the DSCB is the same as the 44-byte data set name you provide. CVAF does not execute the read or write unless they match. If they do not match, CVAF ignores the specified BFLEARG and reads or writes the DSCB according to the rules given in "Specifying a Data Set Name to Read or Write a DSCB" on page 30.

If the buffer is for a 140-byte read or write, CVAF issues a channel program to read or write the DSCB at the location specified in the buffer list entry. CVAF does not use the data set name you specified. If you specify VERIFY=YES, CVAF verifies that the designated DSCB is a format-0 DSCB before issuing the write channel program.

READING OR WRITING VTOC INDEX RECORDS: VIRs may be read or written explicitly using the BUFLIST keyword or may be read implicitly using the IXRCDS and MAPRCDS keywords. A buffer list address may be supplied in the BUFLIST keyword to read or write one or more VIRs. The buffer list header must have the VIR bit set to one and the DSCB bit set to zero. Each entry in the buffer list (and subsequent buffer lists if more than one is chained) is inspected. If the skip bit is set to zero, the RBA bit is set to one (and the CCHHR and TTR bits are set to zero), and the buffer address is nonzero, the entry will be processed. The RBA in the argument field of the buffer list entry is used to read or write a VIR using the buffer address. Read and write requests will be in the order of entries in the buffer list.

Each of the storage areas you provide must be in your program's protect key.

For a write request, the modification bit in the buffer list entries is inspected. If the bit is not set in any entry, all are written. The modification bit is set to zero for entries whose VIR is written.

Map records and the first high-level VTOC index entry record may be read by supplying the keywords MAPRCDS=YES and/or IXRCDS=KEEP and, at the same time, not supplying an address in the CVAF parameter list CVMRCDS/CVIRCDS fields.

READING MAP RECORDS AND VIERS: To read and retain in virtual storage the VTOC index map records and first high-level VIER, either ACCESS=READ or ACCESS=WRITE must be coded. Neither the DSN field nor the BUFLIST field is required.

MAPRCDS=YES must be coded to read and retain map records. The CVAF parameter list field CVMRCDS must be zero. CVAF will obtain a buffer list with the number of entries and buffers required to read all the map VIRs. The buffer list address will be put in the CVMRCDS field by CVAF.

IXRCDS=KEEP is coded in order to read and retain the first high-level VIER and (if an index search is required) all VIERS read. If the CVAF parameter list field CVIRCDS is zero, CVAF will obtain a buffer list with entries and buffers and read the first high-level VIER. The number of entries and number of buffers are determined by CVAF. If CVIRCDS is not zero, only VIERS required for an index search will be read.

The integrity of the maps and VIER read can only be ensured if you are enqueued on the VTOC and, in the case of shared DASD, reserved to the unit.

Map and VIER buffers obtained by CVAF (and retained) must be released by a subsequent CVAF call.

RELEASING BUFFERS AND BUFFER LISTS OBTAINED BY CVAF: There are three ways to release buffers and buffer lists obtained by CVAF.

- Code MAPRCDS=NO or MAPRCDS=(NO,addr) for any specification of ACCESS in order to free the MAP records buffer list.
- Code IXRCDS=NOKEEP or IXRCDS=(NOKEEP,addr) for any specification of ACCESS in order to free the index records buffer list.
- Code ACCESS=RLSE and supply a buffer list address through the BUFLIST keyword for a subsequent CVAF call.

CVAF will free all eligible buffers and any buffer lists if they become empty. Eligible buffers are those pointed to by buffer list entries with the skip bit off. A buffer list will be freed if no buffer list entry has the skip bit on. If buffer lists are chained together, all buffer lists will be checked and freed if appropriate.

You must ensure that you do not request CVAF to release the same buffer list twice by supplying its address in more than one place.

| Accessing DSNs or DSCBs in Sequential Order

Each CVAFSEQ call may request the return of one of the following:

- One format-1 or format-4 DSCB in indexed (data-set-name) order
- One or more DSCBs in physical-sequential order (but only one DSCB can be requested by an unauthorized caller)
- The next data set name in the index

CVAF reads the DSCBs into buffers supplied through the BUFLIST keyword. "CVAFSEQ Macro" on page 190 discusses detailed information about the CVAFSEQ VTOC access macro.

The argument of each DSCB read is also supplied in the buffer list. DSCBs of 96 bytes must be requested in the buffer list for indexed access; 140 bytes is required for physical-sequential access.

If indexed order is chosen, the VTOC index is used to return each format-1 or format-4 DSCB whose name is in the index. An option (DSNONLY=YES) allows only the data set names in the VTOC index, but not the DSCBs, to be obtained. In this case, the CCHHR of the DSCB is returned in the argument area supplied through the ARG keyword. The DSN area supplied is updated at each CVAFSEQ call to contain the data set name of each DSCB read.

INITIATING INDEXED ACCESS (DSN ORDER): To initiate indexed access (DSN order), either supply in the area coded through the DSN keyword 44 bytes of binary zeros (to indicate the first data set name in the index) or supply the data set name you want to serve as the starting place for the index search.

The name returned in the DSN area will be the one equal to or greater than the DSN supplied, depending on the specification of the ACCESS keyword. The DSN field is updated by CVAF.

The ACCESS keyword determines whether the search is for a DSN greater than or equal to that supplied.

If DSNONLY=NO is coded, the DSCB and argument are returned to you, using the buffer list provided through the BUFLIST keyword. The first entry in the buffer list with a skip bit of zero and a nonzero buffer address is used. The argument value is supplied if either the TTR or CCHHR bit is set in the buffer list entry. The default is CCHHR. The DSCB size in the buffer list entry must be 96 bytes for indexed access.

If DSNONLY=YES is coded, the CCHHR argument is supplied in the ARG area.

Note that the data set name of the format-4 DSCB is in the index and that its name (44 bytes of X'04') may be returned to you. The format-4 DSCB's name is likely to be the first data set name in the VTOC index.

INITIATING PHYSICAL-SEQUENTIAL ACCESS: To initiate physical-sequential access, the DSN keyword must be omitted or DSN=0 must be coded. The argument field in the first buffer list entry must be initialized to zero or to the argument of the DSCB to begin the read. If the argument is zero, the argument used will be the start of the VTOC.

The DSCB size must be set to 140 in buffer list entries.

The specification of ACCESS will determine whether the DSCB whose argument is supplied or the DSCB following it is to be read.

For example, to read the first DSCB (the format-4 DSCB) in the VTOC, the BFLEARG in the first buffer list entry may be set to zero, and ACCESS=GTEQ coded in the CVAFSEQ macro. If ACCESS=GT is subsequently coded, the second DSCB (the first format-5 DSCB) is read.

If you are authorized, as many DSCBs as there are entries in the buffer list will be read with a single CVAF call. Only one DSCB will be read if you aren't authorized.

Only one buffer list is used; a second buffer list chained to the first will not be inspected. All entries in the buffer list will be used for authorized callers. The skip bit will not be inspected. Each entry must have a buffer address, the length field set to 140, and the TTR or CCHHR bit set (if neither bit is set, the CCHHR bit will be set on). Only the first entry will be used for unauthorized callers. The argument field of each buffer list entry will be updated by CVAF with the argument of the DSCB. The argument value is returned in either TTR or CCHHR format, depending on whether the TTR or CCHHR bit is set in the buffer list entry. The default is CCHHR.

Only the argument in the first entry is used to begin the search. Arguments in subsequent entries are not inspected. If a nonzero argument value is supplied in the first entry, there must be a DSCB with that argument.

End-of-data is indicated with a return code of 4 in register 15 and CVSTAT set to X'20'. Each buffer list entry following the last DSCB read has its argument field set to zero (this may be the first entry if no DSCBs are read).

Note that all DSCBs, including format-0 DSCBs, are read. You cannot be certain that you have read all format-1 through -6 DSCBs until the entire VTOC has been read. For a nonindexed VTOC, the CCHHR of the last format-1 DSCB is contained in the format-4 DSCB field DS4HPCHR; format-2 through -6 DSCBs may reside beyond that location. For an indexed VTOC, the VMDS contains information about which DSCBs are format-0 DSCBs.

Obtaining Information from the VTOC Index

ACCESS=MAPDATA is used to obtain information contained in the space maps. "CVAFDSM Macro" on page 185 discusses detailed information about the CVAFDSM VTOC access macro.

To count the number of unallocated VIRs in the VTOC index space map (VIXM), COUNT=YES and MAP=INDEX are coded. The number of unallocated VIRs is returned in the 4-byte area supplied through the CTAREA keyword.

To count the number of format-0 DSCBs, COUNT=YES and MAP=VTOC are coded. The number of format-0 DSCBs in the VTOC map of DSCBs VMDS is returned in the 4-byte area supplied through the CTAREA keyword.

To obtain one or more free space extents from the VTOC pack space map (VPSM), COUNT=NO and MAP=VOLUME are coded. The extents are returned in the area supplied through the EXTENTS keyword. Each extent is returned in a 5-byte XXYYZ format, the same as for a format-5 DSCB extent, where XX is the relative track address (RTA) of the first track of the extent, YY is the number of whole cylinders in the extent, and Z is the number of additional tracks in the extent. The RTA supplied to CVAF in the first (or only) extent will serve as a starting point for the VPSM search; the extent returned will be the next free extent with a higher starting RTA than the one supplied.

If all the unallocated extents in the VPSM are supplied before filling in all the extents supplied, the remaining extents are set to zero. Register 15 is set to 4 on return, with the CVSTAT field in the CVPL set to X'20' to indicate the end of data.

DIAGNOSING VTOC ERRORS

Actions Taken When an Error Occurs

These actions are taken if an error occurs:

- If an index structure error is detected, DADSM or CVAF will cause the VTOC index to be disabled. The indexed VTOC bit will be zeroed in the format-4 DSCB. A software error record will be written to SYS1.LOGREC. A system dump is taken. The VTOC will be converted to a nonindexed format at the next DADSM allocate or extend call.
- If a program check, machine check, or other error occurs while using a VTOC access macro, a SYS1.LOGREC message is written, and a system dump is taken.
- An error code is put in the CVSTAT field of the CVPL. The values and explanations of these error codes are listed in Appendix C, "VTOC Index Error Message and Associated Codes" on page 221.

Recovering from System or User Errors

Neither the VTOC nor the VTOC index need be recovered from a user error caused by an unauthorized user, because an unauthorized user cannot modify a VTOC.

A system error will affect a VTOC and VTOC index, probably by interrupting DADSM while it is updating, thus leaving the VTOC and/or the VTOC index in a partially updated state. Both the VTOC and the VTOC index are designed to cause DADSM to recover from such an interruption.

For a nonindexed VTOC (or a VTOC with an index that has been disabled), a subsequent call to DADSM ALLOCATE or EXTEND will

cause VTOC convert routines to reestablish the free space (format-5 DSCBs).

For an indexed VTOC, a subsequent call to any DADSM function will cause the recovery of the previous interrupt (either by backing out or completing the interrupted function).

GTF Trace

A trace function exists to trace all CVAE calls for VTOC index output I/O, all VTOC output I/O, and all VTOC index and space map modifications. For information on this function, see CVAE Diagnosis Reference.

LISTING A VTOC AND VTOC INDEX

A VTOC and VTOC index can be listed using the IEHLIST utility program. Dump, formatted, or abridged listings can be obtained by using the LISTVTOC command of IEHLIST.

CHAPTER 2. EXECUTING YOUR OWN CHANNEL PROGRAMS (EXCP)

The execute-channel-program (EXCP) macro instruction provides you with complete control of the data organization based on device characteristics. This chapter contains a general description of the function and application of the EXCP macro instruction, accompanied by descriptions of specific control blocks and macro instructions used with EXCP. Factors that affect the operation of EXCP, such as device variations and program modification, are also discussed.

Before reading this chapter, you should be familiar with system functions and with the structure of control blocks, as well as with the operational characteristics of the I/O devices required by your channel programs. Operational characteristics of specific I/O devices are described in IBM publications for each device.

You also need to understand the information in these publications:

- Data Administration Guide contains the standard procedures for I/O processing under the operating system.
- Assembler H Version 2 Application Programming: Guide contains the information necessary to code programs in the assembler language.
- Data Administration: Macro Instruction Reference describes the system macro instructions that can be used in programs coded in the assembler language.
- Debugging Handbook, Volumes 2 through 5, contain format and field descriptions of the system control blocks referred to in this chapter.
- Conversion Notebook describes the factors to consider when converting from MVS/370 at the MVS/SP Version 1 level to MVS/XA.

The execute-channel-program (EXCP) macro instruction causes a supervisor-call interruption to pass control to the EXCP processor. (I/O process is the name we will use for the EXCP processor and the I/O supervisor. For our purposes, it's unnecessary to understand how input/output processing is divided between the two.) EXCP also provides the I/O supervisor with control information regarding a channel program to be executed. When an IBM access method is being used, an access method routine is responsible for issuing EXCP. If you are not using an IBM access method, you must issue EXCP in your program. (The EXCP macro instruction cannot be used to process SYSIN or SYSOUT data sets.)

You issue EXCP primarily for I/O programming situations to which the standard access methods do not apply. If you are writing your own access method, you must include EXCP for I/O operations. EXCP must be used for processing nonstandard labels, including reading and writing labels and positioning magnetic tape volumes.

To issue EXCP, you must provide a channel program (a list of channel command words) and several control blocks in your program area. The I/O process then schedules I/O requests for the device you have specified, executes the specified I/O commands, handles I/O interruptions, directs error recovery procedures, and posts the results of the I/O requests.

USING EXCP IN SYSTEM AND PROBLEM PROGRAMS

This section explains the procedures performed by the system and the programmer when EXCP is issued by the routines of IBM access methods. The additional procedures you must perform when issuing EXCP yourself are then described by direct comparison.

HOW THE SYSTEM USES EXCP

When using an IBM access method to perform I/O operations, the programmer is relieved of coding channel programs and constructing the control blocks necessary for the execution of channel programs. To permit I/O operations to be handled by an access method, the programmer need only issue the following macro instructions:

- A DCB macro instruction that produces a data control block (DCB) for the data set to be retrieved or stored
- An OPEN macro instruction that initializes the data control block and produces a data extent block (DEB) for the data set
- A macro instruction (for example, GET or WRITE) that requests I/O operations

Access method routines will then:

1. Create a channel program that contains channel commands for the I/O operations on the appropriate device
2. Construct an input/output block (IOB) that contains information about the channel program
3. Construct an event control block (ECB) that is later posted with a completion code each time the channel program terminates
4. Issue an EXCP macro instruction to pass the address of the IOB to the routines that initiate and supervise the I/O operations

The I/O process consists of:

5. Constructing a request queue element (RQE) for scheduling the request
6. If the requestor is in a V=V address space, fixing the buffers so that they cannot be paged out and translating the requestor's virtual channel program into a real channel program
7. Issuing a start subchannel (SSCH) instruction to cause the channel to execute the real channel program
8. Processing I/O interruptions and scheduling error recovery procedures when necessary
9. Posting a completion code in the event control block after the channel program has been executed

Note: If the requestor is an authorized program in a V=R address space, a real channel program is provided; thus, item 6 is not performed.

The programmer is not concerned with these procedures and does not know the status of I/O operations until they are completed. Device-dependent operations are limited to those provided by the macro instructions of the particular access method selected.

HOW TO USE EXCP IN PROBLEM PROGRAMS

To issue the EXCP macro instruction directly, you must follow the procedures that the access methods would perform, as summarized in items 1 through 4 of the preceding discussion. In addition to constructing and opening the data control block with the DCB and OPEN macro instructions, you must construct a channel program, an input/output block, and an event control block before you can issue EXCP. The I/O process generally handles items 5 through 9.

After issuing EXCP, you should issue a WAIT macro instruction, specifying the address of the event control block, to determine whether the channel program has terminated. If volume switching is necessary, you must issue an EOVS macro instruction. When all processing of the data set has been completed, you must issue a CLOSE macro instruction to restore the data control block.

All external interfaces for EXCP are compatible between MVS/370 and MVS/XA, except for the restrictions noted below. These restrictions relate only to the support of virtual and real addresses above 16-megabyte.

EXCP will be available to programs executing in either 24-bit or 31-bit addressing mode. However, in order to maintain the required compatibility, the following restrictions apply:

- EXCP will only support a 24-bit virtual storage interface. In addition, all areas related to I/O operations (for example, I/O buffers, channel command words, IOBs, DEBs, appendages, and so forth), must remain 24-bit virtual addressable. EXCP (channel command word translator) will allow 24-bit virtual I/O buffers to be fixed above 16-megabyte real. When a channel command word (CCW) references a real address above 16-megabyte, the CCW translator will build an indirect addressing word (IDAW) for that CCW. Note that this is not supported for format-1 CCWs. All virtual addresses must be below 16-megabyte. For V=R users, CCWs and IDAWs must be below 16-megabyte real.
- Only format-0 CCWs are accepted as input.
- All user-specified appendage routines will be given control in 24-bit addressing mode and must return in the same mode.

Note: Access methods run in 24-bit addressing mode. Users running in 31-bit mode must interface to the access methods by using a user-written routine that is resident below 16-megabyte virtual (because the access methods will be able to return control only to a 24-bit addressable location). All addresses, buffers, parameters, control blocks, save areas and exit addresses must be below 16-megabyte virtual. All access methods (except VSAM), for example, GET or PUT, must be called in 24-bit addressing mode.

31-BIT IDAW PROGRAMMING NOTES

A virtual channel program provided by the EXCP caller may have one or more CCWs with the IDA flag set and the address portion of these CCWs pointing to a single 4-byte IDAW. This EXCP function is referred to as virtual IDAWs.

The 4-byte IDAW can contain a virtual address that ranges from 0 to the maximum 31-bit address. Virtual IDAWs are supported on all virtual CCWs except:

- Transfer in channel (TIC) commands.

- All non data-transfer type commands: for example, recalibrate, rewind, set space, fold, block data check, no operation, control commands.
- Read, read backward, and sense commands, with the skip flag set.

The same addressing restrictions apply to EXCPVR inputs with the exception that 31-bit real data areas may be specified by the user-created CCWs through the use of IDAWs. All CCWs and IDAWs must be below 16-megabyte real.

Only format-0 CCWs are accepted as input.

All other areas related to the EXCP/EXCPVR I/O operation (for example, CCWs, IDAWs, IOBs, DEBs, DCBs, appendages, and so forth) must remain 24-bit addressable.

Note, however, that the EXCP processor will allow both 24-bit and 31-bit virtual I/O buffers to be fixed above 16-megabyte real.

HOW EXCP OPERATES IN A V=R ADDRESS SPACE

User-constructed channel programs for I/O operations of an authorized program in a V=R address space are not translated. Because the address space is V=R, any CCWs created by the user have correct real data addresses. (Translation would only re-create the user's channel program, so the CCWs are used directly.)

Modification of an active channel program by data read in or by processor instructions is legitimate in a V=R address space, but not in a V=V address space.

EXCP REQUIREMENTS

This section describes the channel program that you must provide in order to issue EXCP. The control blocks that you must either construct directly or cause to be constructed by use of macro instructions are also described.

All areas related to the EXCP/EXCPVR I/O operation (for example, CCWs, IDAWs, IOBs, DEBs, DCBs, appendages, and so forth) must remain 24-bit addressable.

Note, however, that the EXCP processor will allow both 24-bit and 31-bit virtual I/O buffers to be fixed above 16-megabyte real.

CHANNEL PROGRAM

The channel program supplied by you and executed through EXCP is composed of CCWs on doubleword boundaries. Each channel command word specifies a command to be executed and, for commands initiating data transfer, the area to or from which the data is to be transferred.

Channel command word operation codes used with specific I/O devices can be found in IBM publications for those devices. All channel command word operation codes described in these publications can be used. In addition, both data chaining and command chaining may be used.

To specify either data chaining or command chaining, you must set appropriate bits in the channel command word and indicate the type of chaining in the input/output block. Both data and command chaining should not be specified in the same channel command word; if they are, data chaining takes precedence.

EXCP does not support channel programs that modify themselves, regardless of the method of modification: data chaining, command chaining, or a program to do the modification. The intended modification in virtual storage has no effect on the running real-channel program (see "Modification of a Channel Program during Execution" on page 42).

CONTROL BLOCKS

When using EXCP, you must be familiar with the function and structure of the IOB, the ECB, the DCB, the DEB, and the IDAW. IOB and ECB fields are illustrated under "Control Block Fields" on page 67. DCB fields are illustrated under "Macro Specifications for Use with EXCP" on page 51. The handling of IDAWs is described under "SIO Appendage" on page 72. Descriptions of these control blocks follow.

Input/Output Block (IOB)

The input/output block is used for communication between the problem program and the system. It provides the addresses of other control blocks, and maintains information about the channel program, such as the type of chaining and the progress of I/O operations. You must define the input/output block and specify its address as the only parameter of the EXCP macro instruction.

Event Control Block (ECB)

The event control block provides you with a completion code that describes whether the channel program was completed with or without error. A WAIT macro instruction, which can be used to synchronize I/O operations with the problem program, must identify the event control block. You must define the event control block and specify its address in the input/output block.

Data Control Block (DCB)

The data control block provides the system with information about the characteristics and processing requirements of a data set to be read or written by the channel program. A data control block must be produced by a DCB macro instruction that includes parameters for EXCP. If appendages are not being used, a short DCB is constructed. Such a DCB does not support reduced error recovery. You specify the address of the data control block in the input/output block.

All DCBs must be located in storage that is not fetch-protected, or, if the task is authorized, in storage that is in the key of the task (TCB KEY).

Data Extent Block (DEB)

The data extent block contains one or more extent entries for the associated data set and other control information. An extent defines all or part of the physical boundaries on an I/O device occupied by, or reserved for, a particular data set. Each extent entry contains the address of a unit control block (UCB) that provides information about the type and location of an I/O device. More than one extent entry can contain the same UCB address. For all I/O devices supported by the operating system, the data extent block is produced during execution of the OPEN macro instruction for the data control block. The system places the address of the data extent block into the data control block. All DEBs must be located in storage that is not fetch-protected, or, if the task is authorized, in storage that is in the key of the task (TCB key). Only authorized tasks (APF-authorized or TCB PKF=0-7) may build DEBs to be used for I/O operations.

HOW THE CHANNEL PROGRAM EXECUTES

This section explains how the system uses your channel program and control blocks after you issue EXCP.

INITIATION OF THE CHANNEL PROGRAM

By issuing EXCP, you request the execution of the channel program specified in the input/output block. The I/O process validates the request by checking certain fields of the control blocks associated with this request. If the I/O process detects invalid information in a control block, it initiates abnormal termination procedures.

The EXCP processor gets:

- The address of the data control block from the input/output block
- The address of the data extent block from the data control block
- The address of the unit control block from the data extent block

It places the IOB, TCB, DEB, and UCB addresses and other information about the channel program into an area called a request queue element (RQE). (Unless you are providing appendage routines (described under "Appendages" on page 43) you should not be concerned with the contents of RQEs.)

If you have provided a start I/O (SIO) appendage, the EXCP processor now passes control to it. The return address from the SIO appendage determines whether the EXCP processor must:

- Execute the I/O operation normally, or
- Skip the I/O operation.

For a description of the SIO appendage and its linkage to the EXCP processor, see "Appendages" on page 43.

If you are issuing EXCP from a V=V address space, the channel program you construct contains virtual addresses. Because channel subsystems cannot use virtual addresses, the EXCP processor must:

- Translate your virtual channel program into one that uses only real addresses.
- Fix in real storage the pages used as I/O areas for the data transfer operations specified in your channel program.

The EXCP processor builds the translated (real) channel program in a portion of real storage.

For direct access devices, specify the seek address in the input/output block. The I/O supervisor constructs a command chain to issue the seek and the set file mask specified in the data extent block, and to pass control to your real channel program.

If your channel program begins with a locate-record command, the I/O process builds a define-extent command and passes control to your real channel program. (You cannot issue the initial seek, set file mask, or define extent. The file mask is set to prohibit seek-cylinder commands, or, if space is allocated by tracks, seek-head commands. If the data set is open for INPUT, write commands are also prohibited.)

For a magnetic tape device, the I/O supervisor constructs a command chain to set the mode specified in the data extent block

and passes control to your real channel program. (You cannot set the mode yourself.)

If the I/O device is other than a direct access device or a magnetic tape device, the I/O supervisor then places the starting CCW of the channel program into the operation request block (ORB) and issues a start subchannel (SSCH) instruction.

MODIFICATION OF A CHANNEL PROGRAM DURING EXECUTION

Any problem program that modifies an active channel program with processor instructions or with data read in by an I/O operation must be run in a V=R address space. It cannot run in a V=V address space because of the channel program translation performed by the I/O supervisor. (In a V=V address space, an attempt to modify an active channel program affects only the virtual image of the channel program, not the real channel program being executed by the channel subsystem.)

A program of this type can be changed to run in a V=V address space by issuing another EXCP macro for the modified portion of the channel program.

COMPLETION OF EXECUTION

The system considers the channel program completed when it receives an indication of a channel-end condition in the subchannel status word (SCSW). Unless a channel-end or abnormal-end appendage directs otherwise, the request queue element for the channel program is made available, and a completion code is placed into the event control block. The completion code indicates whether errors are associated with channel end. If device end occurs simultaneously with channel end, errors associated with device end (that is, unit exception or unit check) are also accounted for.

If device end occurs after channel end and if an error is associated with device end, the completion code in the event control block does not indicate the error. However, the status of the unit and channel is saved by the I/O supervisor for the device, and the UCB is marked as intercepted. The input/output block for the next request directed to the I/O device is also marked as intercepted. The error is assumed to be permanent, and the completion code in the event control block for the intercepted request indicates interception. The DCBIFLGS field of the data control block is also flagged to indicate a permanent error. Note that, if a write-tape-mark or erase-long-gap CCW is the last or only CCW in your channel program, the I/O process will not attempt recovery procedures for device end errors. In these circumstances, command chaining a NOP CCW to your write-tape-mark or erase-long-gap CCW ensures initiation of device-end error recovery procedures.

To be prepared for device-end errors, you should be familiar with device characteristics that can cause such errors. After one of your channel programs has terminated, you should not release buffer space until you have determined that your next request for the device has not been intercepted. You may reissue an intercepted request.

INTERRUPTION HANDLING AND ERROR RECOVERY PROCEDURES

An I/O interruption allows the processor to respond to signals from an I/O device that indicate either termination of a phase of I/O operations or external action on the device. A complete explanation of I/O interruptions is contained in IBM System/370 XA Principles of Operation. For descriptions of interruption by specific devices, see the IBM publications for each device.

If error conditions are associated with an interruption, the I/O supervisor schedules the appropriate device-dependent error

routine. The channel subsystem is then restarted with another request that is not related to the channel program in error. (The following paragraphs discuss "related" channel programs.) If the error recovery procedures fail to correct the error, the system places ones in the first two bit positions of the DCBIFLGS field of the data control block. You are informed of the error by an error code in the event control block.

If a channel program depends on the successful completion of a previous channel program (as when one channel program retrieves data to be used in building another), the previous channel program is called a "related" request. Such a request must be identified to the EXCP processor. To find out how to do this, see "Input/Output Block (IOB) Fields" on page 67.

If a permanent error occurs in the channel program of a related request, the EXCP processor removes the request queue elements for all dependent channel programs from their queue and makes them available.

The related request queue (RRQ) reflects the order in which request queue elements are removed from their queue.

For all requests dependent on the channel program in error, the system places completion codes into the event control blocks. The DCBIFLGS field of the data control block is also flagged. Any requests for a data control block with error flags are posted complete without execution. To reissue requests dependent on the channel program in error, you must reset the first two bits of the DCBIFLGS field of the data control block to zeros. You then reissue EXCP for each channel program desired.

With the IBM 3800 Printing Subsystem, a cancel key or a system-restart-required paper jam causes both a lost data indicator to be set in DCBIFLGS and a lost page count and channel page identifier to be stored in the UCB extension. (See JES3 Data Areas, TSO/E Data Areas, and IBM 3800 Printing Subsystem Programmer's Guide.)

APPENDAGES

An appendage is a programmer-written routine that provides additional control over I/O operations. By using appendages, you can examine the status of I/O operations and determine the actions to be taken for various conditions. An appendage may receive control when one of the following occurs:

- EXCP SVC
- Program-controlled interrupt
- End of extent
- Channel end
- Abnormal end

Appendages get control in supervisor state, receiving the following pointers from the EXCP processor:

- Register 1: Points to the request queue element for the channel program.
- Register 2: Points to the input/output block (IOB).
- Register 3: Points to the data extent block (DEB).
- Register 4: Points to the data control block (DCB).
- Register 6: Points to the seek address if control is given to an end-of-extent appendage.

- Register 7: Points to the unit control block (UCB).
- Register 13: Points to a 16-word area you can use to save input registers or data.
- Register 14: Points to the location in the EXCP processor where control is to be returned after execution of an appendage. When returning control to the EXCP processor, you may use displacements from the return address in register 14. Allowable displacements are summarized in Figure 12 and described later for each appendage.
- Register 15: Points to the entry point of the appendage.

The processing done by appendages is subject to these requirements and restrictions:

- Register 9, if used, must be set to binary zeros before control is returned to the system. All other registers, except those indicated in the descriptions of each appendage, must be saved and restored if they are used. Figure 12 summarizes register conventions.
- No SVC instructions or instructions that change the status of the system (for example, WTO, LPSW, or any privileged instructions) can be issued.
- Loops that test for the completion of I/O operations must not be used.
- Storage used by the I/O supervisor or EXCP processor must not be altered.

The types of appendages are described in the following sections, with explanations of when they are created, how they return control to the system, and which registers they may use without saving and restoring their contents.

Note: All user-specified appendages will be given control in 24-bit addressing mode and must return in the same mode.

Appendage	Entry Point	Returns	Available Work Registers ¹
EOE	Reg 15	Reg 14 + 0 Reg 14 + 4 Reg 14 + 8	Return Skip Try Again Reg. 10, 11, 12, and 13
SIO	Reg 15	Reg 14 + 0 Reg 14 + 4	Normal Skip Reg. 10, 11, and 13
PCI	Reg 15	Reg 14 + 0	Normal Reg. 10, 11, 12, and 13
CHE	Reg 15	Reg 14 + 0 Reg 14 + 4 Reg 14 + 8 Reg 14 + 12	Normal Skip Re-EXCP By-Pass Reg. 10, 11, 12, and 13
ABE	Reg 15	Reg 14 + 0 Reg 14 + 4 Reg 14 + 8 Reg 14 + 12	Normal Skip Re-EXCP By-Pass Reg. 10, 11, 12, and 13

Figure 12. Entry Points, Returns, and Available Work Registers for Appendages

Note to Figure 12:

¹ Certain register conventions for passing parameters from appendages to the EXCP processor must be followed. These conventions are described in the individual appendage descriptions.

START-I/O (SIO) APPENDAGE

Unless an error recovery procedure is in control, the EXCP processor passes control to the SIO appendage just before the EXCP processor translates your channel program.

Optional return vectors give the I/O requestor the following choices:

Reg. 14 + 0

Normal return. Normal channel program translation and initiation of I/O.

Reg. 14 + 4

Skip the I/O operation. The channel program is not posted complete, but the request queue element is made available. You may post the channel program as follows:

1. Save necessary registers.
2. Put the address of the post routine (found at CVTOPT01 in the communications vector table) in register 15.
3. Place the ECB address from the IOB in register 11.
4. Set the completion code in register 10. These are the four bytes of an ECB.
5. Go to the post routine pointed to by the CVT, using BALR 14, 15.

PROGRAM-CONTROLLED INTERRUPTION (PCI) APPENDAGE

This appendage is entered at least once if the channel finds one or more PCI bits on in a channel program. It may be entered as many times as the channel finds PCI bits on. Before the appendage is entered, the contents of the subchannel status word are placed in the "channel status word" field of the input/output block.

A PCI appendage will be reentered if an error recovery procedure is retrying a channel program in which a PCI bit is on. The IOB error flag is set when the error recovery procedure is in control (IOBFLAG1 = X'20'). (For special PCI conditions encountered with command retry, see "Channel Programming Considerations" on page 50.)

To post the channel program from a PCI appendage, the procedure described for the start-I/O appendage is used if the step is running ADDRSPC=VIRT or an authorized program is running V=R. If the step is running ADDRSPC=REAL and an authorized program issued the EXCP request or if SVC 114(EXCPVR) was issued, the PCI appendage uses real storage addresses, and the following procedure is used to post the channel program from the PCI appendage.

1. Put the completion code in register 10 and place X'80' in the high-order byte to indicate the key is in register 0 (step 5).
2. Put X'80' in the high-order byte of register 11 and the address of the ECB in the low-order bytes.
3. Put X'80' in the high-order byte of register 12 and the address of a BR 14 instruction in the low-order bytes. This BR 14 must be in storage addressable from any address space (for example, CVTBRET) and must be in storage addressable by 24 bits. Note that only registers 9 and 14 are restored when you use this option.

- Put the address of the ASCB in register 13.

The next two paragraphs describe how to obtain the ASCB address and are followed by sample instructions to illustrate the procedure.

Get the SRB address associated with the I/O operation from the RQE field, RQESRB (the RQE address was in register 1 when the appendage was given control). Get the IOSB address from SRBPARM. From that IOSB, get the identifier field, IOSASID. Multiply IOSASID by 4.

Get the pointer to the ASVT (address space vector table) found at CVTASVT. The address of the ASCB can be found in the ASVT, using the field ASVTENTY-4 indexed by the value calculated in the above paragraph.

```
USING    RQE,1
L        Y,RQESRB
USING    SRBSECT,Y
LH       Y,SRBPARM
USING    IOSB,Y
LH       Y,IOSASID
SLA     Y,2
L        X,16
USING    CVT,X
L        X,CVTASVT
USING    ASVT,X
L        13,ASVTENTY-4(Y)
```

Note: X and Y are work registers.

- Put the requestor's key in register 0.
- Put the address of the post routine (found at CVTOPT01 in the communications vector table) in register 15.
- Go to the post routine using BALR 14,15. Upon return, only registers 9 and 14 are valid. For more information on the POST routine, see Supervisor Services and Macro Instructions.

This procedure can be used even if the PCI appendage uses virtual storage addresses, but performance may be slightly slower.

To return control to the EXCP processor for normal interruption processing, use the return address in register 14.

END-OF-EXTENT (EOE) APPENDAGE

This appendage is entered when the seek address specified in the input/output block is outside the allocated extent limits indicated in the data extent block.

If you use the return address in register 14 to return control to the system, the abnormal-end appendage is entered. An end-of-extent error code (X'42') is placed in the "ECB code" field of the input/output block for subsequent posting in the ECB.

You may use the following optional return addresses:

- Contents of register 14 plus 4: The channel program is posted complete; its request element is returned to the available queue.
- Contents of register 14 plus 8: The request is tried again.

You may use registers 10 through 13 in an end-of-extent appendage without saving and restoring their contents.

Note: If an end-of-cylinder or file-protect condition occurs, the EXCP processor updates the seek address to the next higher cylinder or track address and reexecutes the request. If the new seek address is within the data set's extent, the request is executed; if the new seek address is not within the data set's extent, the end-of-extent appendage is entered. If you want to try the request in the next extent, you must move the new seek address to the location pointed to by register 6.

If a file protect is caused by a full seek (command code=07) embedded within a channel program, the request is flagged as a permanent error, and the abnormal end appendage is entered.

CHANNEL-END (CHE) APPENDAGE

This appendage is entered when a channel end (CHE), unit exception (UEX) with or without channel end or when channel end with wrong-length record (WLR) occurs without any other abnormal-end conditions.

If you use the return address in register 14 to return control to the EXCP processor, the channel program is posted complete, and its request element is made available. In the case of unit exception or wrong-length record, the error recovery procedure is performed before the channel program is posted complete, and the IOBIOERR flag (X'04') in IOBFLAG1 is set on. The CSW status may be obtained from the IOBCSW field.

If the appendage takes care of the wrong-length record or unit exception or both, it may turn off the IOBIOERR flag in IOBFLAG1 and return normally. The event will then be posted as complete (completion code X'7F' under normal conditions, taken from the high-order byte of the IOBECBCC field). If the appendage returns normally without resetting the IOBIOERR flag to zero, the request will be routed to the associated device error recovery procedure (ERP), and the abnormal-end appendage will then be entered with the completion code in IOBECBCC set to X'41' if the ERP could not correct the error. (See Step 1 of "Abnormal-End (ABE) Appendage" on page 48.)

You may use the following optional return addresses:

- Contents of register 14 plus 4: The channel program is not posted complete, but its request element is made available. You may post the channel program by using the calling sequence described under the start-I/O appendage. This is especially useful if you want to post an ECB other than the ECB in the input/output block.
- Contents of register 14 plus 8: The channel program is not posted complete, and its request element is placed back on the request queue so that the I/O operation can be retried. For correct reexecution of the channel program, you must reinitialize the IOBFLAG1, IOBFLAG2, and IOBFLAG3 fields of the input/output block and set the "Error Counts" field to zero. As an added precaution, the IOBSENS0, IOBSENS1, and IOBCSW fields should be cleared.
- Contents of register 14 plus 12: The channel program is not posted complete, and its request element is not made available. (This return must be used if, and only if, the appendage has passed the RQE to the exit effector for use in scheduling an asynchronous routine. For information on the exit effector, see System Macros and Facilities.)

You may use registers 10 through 13 in a channel-end appendage without saving and restoring their contents.

ABNORMAL-END (ABE) APPENDAGE

This appendage may be entered on abnormal conditions, such as unit check, unit exception, wrong-length indication, program check, protection check, channel data check, channel control check, interface control check, chaining check, out-of-extent error, and intercept condition (that is, device end error). It may also be entered when an EXCP is issued for a request queue element that has already been purged.

1. When this appendage is entered because of a unit exception or wrong-length record indication or both, IOBECBCC is set to X'41'. For further information on these conditions, see "Channel-End (CHE) Appendage" on page 47.
2. When the appendage is entered because of an out-of-extent error, the IOBECBCC is set to X'42'.
3. When this appendage is entered with IOBECBCC set to X'4B', it is because of:
 - a. The tape error recovery procedure (ERP) encountering an unexpected load point, or
 - b. The tape error recovery procedure (ERP) finding zeros in the command address field of the CSW.
4. When the appendage is first entered because of an intercept condition, the IOBECBCC is set to X'7E'. If it is then determined that the error condition is permanent, the appendage will be entered a second time with the IOBECBCC set to X'44'. The intercept condition signals that an error was detected at device end after channel end on the previous request.
5. When the appendage is entered because of an EXCP being issued to an already purged request queue element, this request will enter the abnormal end appendage with the IOBECBCC set to X'48'. This applies only to related requests.
6. If the appendage is entered with IOBECBCC set to X'7F', it may be because of a unit check, program check, protection check, channel data check, channel control check, interface control check, or chaining check. If the IOBECBCC is X'7F', it is the first detection of an error in the associated channel program. If the IOBEX (X'04") flag (bit 5 of the IOBFLAG1) is on, the IOBECBCC field will contain a X'41', X'42', X'48', X'4B', or X'4F' in hexadecimal, indicating a permanent I/O error.

To determine if an error is permanent, you should check the IOBECBCC field of the IOB. To determine the type of error, check the channel status word field and the sense information in the IOB. However, when the IOBECBCC is X'42', X'48', or X'4F', these fields are not applicable. For X'44', the CSW is applicable, but the sense is valid only if the unit check bit is set.

If you use the return address in register 14 to return control to the system, the channel program is posted complete, and its request element is made available. You may use the following optional return addresses:

- Contents of register 14 plus 4: The channel program is not posted complete, but its request element is made available. You may post the channel program by using the calling sequence described under the start-I/O appendage.
- Contents of register 14 plus 8: The channel program is not posted complete, and its request element is placed back on the request queue so that the request can be retried. Reinitialize the IOBFLAG1, IOBFLAG2, and IOBFLAG3 fields of the input/output block and set the IOBERRCT field to zero.

As an added precaution, clear the IOBSENS0, IOBSENS1, and IOBCSW fields.

- Contents of register 14 plus 12: The channel program is not posted complete, and its request element is not made available. (This return must be used if, and only if, the appendage has passed the RQE to the exit effector for use in scheduling an asynchronous routine.)

You may use registers 10 through 13 in an abnormal-end appendage without saving and restoring their contents.

MAKING YOUR APPENDAGES PART OF THE SYSTEM

Before your appendages can be executed, they must become members of either the SYS1.LPALIB or SYS1.SVCLIB data set. There are two ways to put appendages into SYS1.LPALIB or SYS1.SVCLIB: They can be included at system generation using the DATASET macro instruction (a full explanation appears in Installation System Generation), or they can be link-edited into SYS1.LPALIB or SYS1.SVCLIB after the system has been generated. Each appendage must have an 8-character member name, the first six characters being IGG019 and the last two being anything in the range from WA to Z9. Note, however, if your program runs in a V=R address space and uses a PCI appendage, the PCI appendage and any appendage that the PCI appendage refers to must be placed in either SYS1.SVCLIB or the fixed link pack area (LPA). For information on providing a list of programs to be fixed in storage, see Initialization and Tuning.

THE AUTHORIZED APPENDAGE LIST (IEAAPP00)

If an "unauthorized" program opens a DCB to be used with an EXCP macro instruction, the names of any appendages associated with the DCB must be listed in the IEAAPP00 member of SYS1.PARMLIB. (An "unauthorized" program is one that runs in a protection key greater than 7 and has not been marked as authorized by the Authorized Program Facility.)

If your appendages were put in SYS1.LPALIB or SYS1.SVCLIB at system generation, their names are automatically put in IEAAPP00. If your appendages were added to SYS1.LPALIB or SYS1.SVCLIB after system generation, you can add IEAAPP00 to SYS1.PARMLIB and put the names of the appendages in it in one job step with the IEBUPDTE utility.

Here is an example of JCL statements and IEBUPDTE input that will add IEAAPP00 to SYS1.PARMLIB and put the names of one EOE appendage, two SIO appendages, two CHE appendages, and one ABE appendage in IEAAPP00:

```
//          JOB          ...
//          EXEC        PGM=IEBUPDTE
//SYSPRINT DD          SYSOUT=A
//SYSUT2   DD          DSN=SYS1.PARMLIB,DISP=OLD
//SYSIN    DD          *
./         ADD         NAME=IEAAPP00,LIST=ALL
EOEAPP WA,
SIOAPP X1,X2,
CHEAPP Z3,Z4,
ABEAPP Z2
/*
```

Note the following about the IEBUPDTE input:

- The type of appendage is identified by six characters that begin in column 1. EOEAPP identifies an EOE appendage, SIOAPP an SIO appendage, CHEAPP a CHE appendage, and ABEAPP an ABE appendage. (The PCI appendage identifier, PCIAPP, is not shown, because the example adds no PCI appendage name to IEAAPP00.)

- Only the last two characters in an appendage's name are specified, beginning in column 8.
- Each statement that identifies one or more appendage names ends in a comma, except the last statement.

You can also use IEBUPDTE to add appendage names later or to delete appendage names. Here is an example of JCL statements and IEBUPDTE input that adds the names of a PCI and an ABE appendage to the IEAAPPO0 appendage list that was created in the preceding example, and deletes the name of an SIO appendage from that list:

```

//          JOB. . .
//          EXEC      PGM=IEBUPDTE
//SYSPRINT DD          SYSOUT=A
//SYSUT2   DD          DSN=SYS1.PARMLIB,DISP=OLD
//SYSIN    DD          *
./         REPL       NAME=IEAAPPO0,LIST=ALL
PCIAPP Y1,
EOEAPP WA,
SIOAPP X1,X2,
CHEAPP Z3,Z4,
ABEAPP Z2,Z4
/*

```

Note the following about the IEBUPDTE input:

- The command to IEBUPDTE in this case is REPL (replace).
- All the appendage names that are to remain in IEAAPPO0 are repeated.
- IGG019Z4 is both a CHE and an ABE appendage.

CHANNEL PROGRAMMING CONSIDERATIONS

Command retry is a function of the channel supporting the IBM 2305-2, 3330/3333, 3340/3344, 3350, 3375, and 3380 direct access devices. When the channel subsystem receives a retry request, it repeats the execution of the CCW, requiring no additional input/output interrupts. For example, a control unit may initiate a retry procedure to recover from a transient error.

A command retry during the execution of a channel program may cause any of the following conditions to be detected by the initiating program:

- **Modifying CCWs:** A CCW used in a channel program must not be modified before the CCW operation has been successfully completed. Without the command retry function, a command was fetched only once from storage by a channel. Therefore, a program could determine through condition codes or program controlled interruptions (PCI) that a CCW had been fetched and accepted by the channel. This permitted the CCW to be modified before reexecution. With the command retry function, this cannot be done, because the channel will fetch the CCW from storage again on a command retry sequence. In the case of data chaining, the channel will retry commands starting with the first CCW in the data chain.
- **Program Controlled Interrupts (PCI):** A CCW containing a PCI flag may cause multiple program-controlled interrupts to occur. This happens if the PCI-flagged CCW was retried during a command retry procedure and if a PCI could be generated each time the CCW is reexecuted.
- **Residual Count:** If a channel program is prematurely terminated during the retry of a command, the residual count in the channel status word (CSW) will not necessarily indicate how much storage was used. For example, if the control unit detects a "wrong-length record" error

condition, an erroneous residual count is stored in the CSW until the command retry is successful. When the retry is successful, the residual in the CSW reflects the correct length of the data transfer.

- Command Address: When data chaining with command retry, the CSW may not indicate how many CCWs have been executed at the time of a PCI. For example:

CCW# Channel Program

1	Read, data chain
2	Read, data chain
3	Read, data chain, PCI
4	Read, command chain

In this example, assume that the control unit signals command retry on Read #3 and the processor accepts the PCI after the channel resets the command address to Read #1 because of command retry. The CSW stored for the PCI will contain the command address of Read #1 when the channel has actually progressed to Read #3.

- Testing Buffer Contents on Data Read: Any program that tests a buffer to determine when a CCW has been executed and continues to execute based on this data may get incorrect results if an error is detected and the CCW is retried.

MACRO SPECIFICATIONS FOR USE WITH EXCP

If you are using the EXCP macro instruction, you must also use DCB, OPEN, CLOSE, and, in some cases, the EOVS macro instruction. The parameters of these macro instructions and the EXCP macro instructions are explained here. A diagram of the data control block is included with the description of the DCB macro instruction.

DEFINING DATA CONTROL BLOCKS FOR EXCP (DCB)

The EXCP form of the DCB macro instruction produces a data control block that can be used with the EXCP macro instruction. You must issue a DCB macro instruction for each data set to be processed by your channel programs. Notation conventions and format illustrations of the DCB macro instruction are given in Data Administration: Macro Instruction Reference. DCB parameters that apply to EXCP may be divided into four categories, depending on the following portions of the data control block that are generated when they are specified:

- Foundation block. This portion is required and is always 12 bytes in length. You must specify two of the parameters in this category.
- EXCP interface. This portion is optional. If you specify any parameter in this category, 20 bytes are generated.
- Foundation block extension and common interface. This portion is optional and is always 20 bytes in length. If this portion is generated, the device-dependent portion is also generated.
- Device dependent. This portion is optional and is generated only if the foundation block extension and common interface portion is generated. Its size ranges from 4 to 20 bytes, depending on specifications in the DEVD parameter. However, if you do not specify the DEVD parameter (and the foundation extension and common interface portion is generated), the maximum 20 bytes for this portion are generated.

Some of the procedures performed by the system when the data control block is opened and closed (such as writing file marks for output data sets on direct access volumes) require

information from optional data control block fields. You should make sure that the data control block is large enough to provide all information necessary for the procedures you want the system to handle.

Figure 13 on page 53 shows the relative position of each portion of an opened data control block. The fields corresponding to each parameter of the DCB macro instruction are also designated, with the exception of DDNAME, which is not included in a data control block that has been opened. The fields identified in parentheses represent system information that is not associated with parameters of the DCB macro instruction.

Sources of information for data control block fields other than the DCB macro instruction are data definition (DD) statements, data set labels, and data control block modification routines. You may use any of these sources to specify DCB parameters. However, if a particular portion of the data control block is not generated by the DCB macro instruction, the system does not accept information intended for that portion from any alternative source.

You may provide symbolic names for the fields in one or more EXCP DCBs by coding a DCBD macro to generate a dummy control section (DSECT). To map the common interface, foundation block extension, and foundation block, you code DSORG=XE. To map the foundation block and EXCP interface, you code DSORG=XA. You may code DSORG=(XA,XE) to map both. For further information, see Data Administration: Macro Instruction Reference.

Foundation Block Parameters

DDNAME=symbol

The name of the data definition (DD) statement that describes the data set to be processed. This parameter must be given.

MACRF=(E)

The EXCP macro instruction is to be used in processing the data set. This operand must be coded.

REPOS={Y|N}

Magnetic tape volumes: This parameter indicates to the dynamic device reconfiguration (DDR) routine whether the user is keeping an accurate block count. If the user is keeping an accurate block count, the DDR routine can attempt to swap the volume. (You must maintain the block count in the DCBBLKCT field.)

Y—The user is keeping an accurate block count, and the DDR routine can attempt to swap the volume.

N—The block count is unreliable, and the DDR routine cannot and will not attempt to swap the volume.

If the operand is omitted, N is assumed.

EXCP Interface Parameters

EOEA=symbol

2-byte identification of an EOE appendage that you have entered into SYS1.LPALIB or SYS1.SVCLIB.

PCIA=symbol

2-byte identification of a PCI appendage that you have entered into SYS1.LPALIB or SYS1.SVCLIB.

SIOA=symbol

2-byte identification of a SIO appendage that you have entered into SYS1.LPALIB or SYS1.SVCLIB.

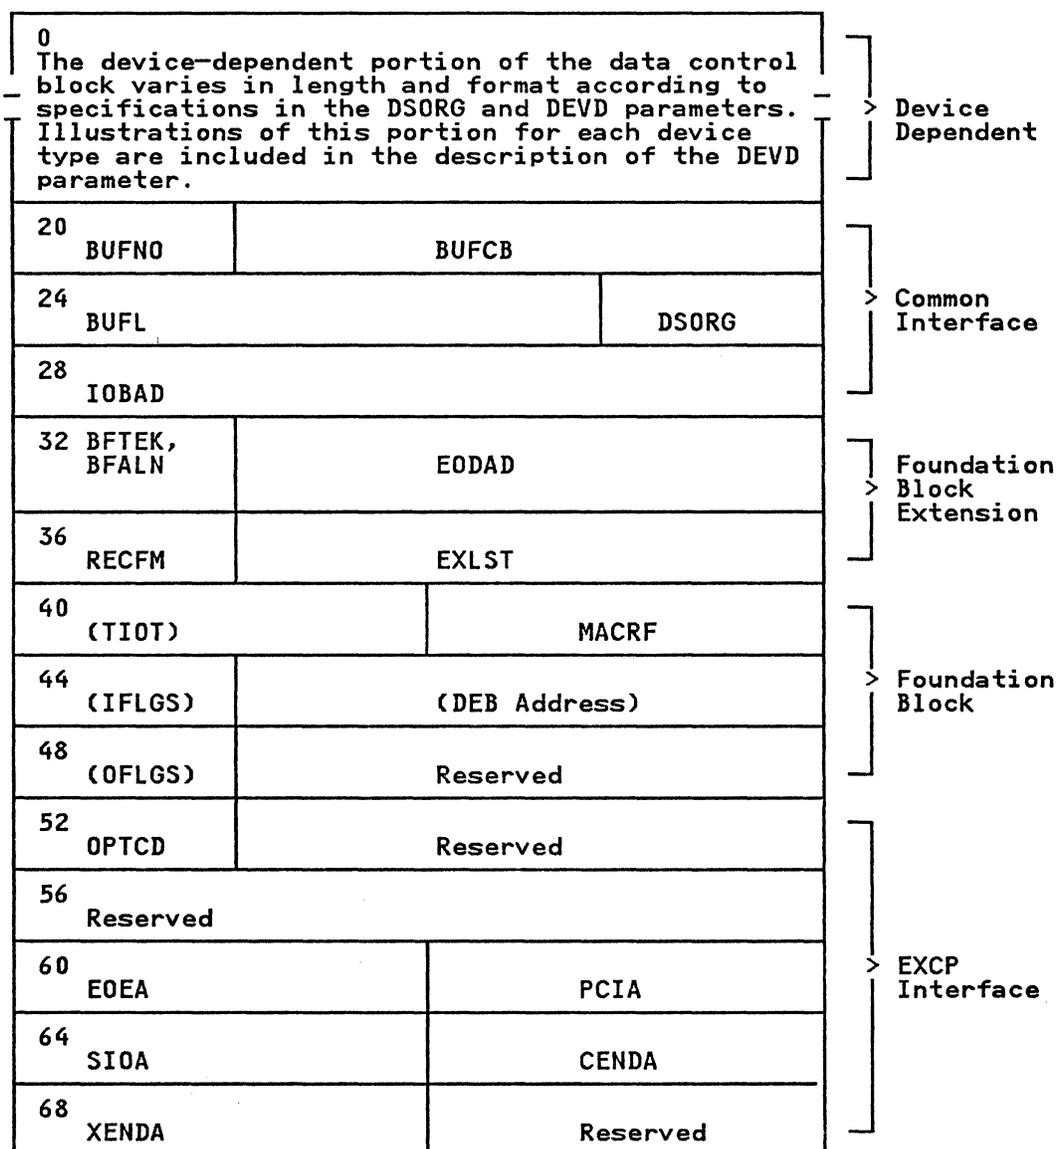


Figure 13. Data Control Block (DCB) Format for EXCP (After OPEN)

CENDA=symbol

2-byte identification of a CHE appendage that you have entered into SYS1.LPALIB or SYS1.SVCLIB.

XENDA=symbol

2-byte identification of an ABE appendage that you have entered into SYS1.LPALIB or SYS1.SVCLIB.

OPTCD=Z

indicates that, for magnetic tape (input only), a reduced error recovery procedure (5 reads only) will occur when a data check is encountered. It should be specified only when the tape is known to contain errors and the application does not require that all records be processed. Its proper use would include error frequency analysis in the SYNAD routine. Specification of this parameter will also cause generation of a foundation block extension.

This parameter is ignored unless it was selected at system generation.

IMSK=value

Any specification indicates that the system will not use IBM-supplied error routines.

Foundation Block Extension and Common Interface Parameters

EXLST=address

the address of an exit list that you have written for exception conditions. The format of this exit list is given in Data Administration Guide.

EODAD=address

the address of your end-of-data-set routine for input data sets. If this routine is not available when it is required, the task is abnormally terminated.

DSORG={PS|PO|DA|IS}

the data set organization (one of the following codes). Each code indicates that the format of the device-dependent portion of the data control block is to be similar to that generated for a particular access method:

Code	DCB Format for
PS	QSAM or BSAM
PO	BPAM
DA	BDAM
IS	QISAM or BISAM

For direct access devices, if you specify PS or PO, you must maintain the following fields of the device-dependent portion of the data control block so that the system can write a file mark for output data sets:

- The track balance (DCBTRBAL) field that contains a 2-byte binary number that indicates the remaining number of bytes on the current track. This number can be obtained from the system track algorithm routine.
- The full disk address (DCBFDAD) field that indicates the location of the current record. The address is in the form MBBCCHHR.

These fields are written into the format-1 DSCB and are used by Open routines for staging MSS data sets. Staging is done only up through the last cylinder specified by these fields if the data set is reopened for OUTPUT, INOUT, OUTIN, OUTINX, or EXTEND.

If you specify PO for a direct access device, the DCBDIRCT field will not be updated. Therefore, you should be careful when using EXCP with the STOW macro.

IOBAD=address

the address of an input/output block (IOB). If a pointer to the current IOB is not required, you may use this field for any purpose.

The following parameters are not used by the EXCP routines. They provide additional information that the system will store for later use by access methods that read or update the data set.

RECFM=code

the record format of the data set. Record format codes are given in Data Administration: Macro Instruction Reference. When writing a data set to be read later, RECFM, LRECL, and BLKSIZE should be specified to identify the data set attributes. LRECL and BLKSIZE can only be specified in a

DD statement, because these fields do not exist in a DCB used by EXCP.

BFTEK={S|E}
the buffer technique, either simple or exchange.

BFALN={F|D}
the word boundary alignment of each buffer, either fullword or doubleword.

BUFL=length
the length in bytes of each buffer; the maximum length is 32767.

BUFNO=number
the number of buffers assigned to the associated data set; the maximum number is 255.

BUFCB=address
the address of a buffer pool control block, that is, the 8-byte field preceding the buffers in a buffer pool.

Device-Dependent Parameters

DEV D=code
the device in which the data set may reside. The codes are listed in order of descending space requirements for the data control block:

Code	Device
DA	Direct access
TA	Magnetic tape
PR	Printer
PC	Card punch
RD	Card reader

Note: For MSS virtual volumes, DA should be used.

If you do not want to select a specific device until job setup time, you should specify the device type requiring the largest area; that is, DEV D=DA.

The following diagrams illustrate the device-dependent portion of the data control block for each combination of device type specified in the DEV D parameter and data set organization specified in the DSORG parameter. Fields that correspond to device-dependent parameters in addition to DEV D are indicated by the parameter name. For special services, you may have to maintain the fields shown in parentheses. The special services are explained in the note that follows the diagram.

Device-dependent portion of data control block when DEVD=DA and DSORG=PS:

4 Reserved	5 DCBFDAD	
8	13 DCBDVTBL	
16 DCBKEYLE	17 DCBDEVT	18 DCBTRBAL

For output data sets, the system uses the contents of the full disk address (DCBFDAD) field, plus one, to write a file mark when the data control block is closed, provided the track balance (DCBTRBAL) field indicates that space is available. If DCBTRBAL is less than 8, the file mark is written on the next sequential track. You must maintain the contents of these two fields yourself if the system is to write a file mark. OPEN will initialize DCBDVTBL and DCBDEVT.

Device-dependent portion of data control block when DEVD=DA and DSORG=DA:

16 DCBKEYLE	18 Reserved
----------------	----------------

Device-dependent portion of data control block when DEVD=TA and DSORG=PS:

12 DCBBLKCT			
16 DCBTRTCH	17 Reserved	18 DCBDEN	19 Reserved

The system uses the contents of the block count (DCBBLKCT) field to write the block count in trailer labels when the data control block is closed or when the EOVS macro instruction is issued. You must maintain the contents of this field yourself if the system is to have the correct block count. (Note: The I/O supervisor increments this field by the contents of the IOBINCAM field of the IOB at the completion of each I/O request.)

When using EXCP to process a tape data set open at a checkpoint, you must be careful to maintain the correct count; otherwise, the system may position the data set incorrectly when restart occurs. If REPOS=Y, the count must be maintained by you for repositioning during dynamic device reconfiguration.

Device-dependent portion of data control block when DEVD=PR and DSORG=PS:

16 DCBPRTSP	18 Reserved
----------------	----------------

Device-dependent portion of data control block when DEVD=PC or RD and DSORG=PS:

16 DCBMODE, DCBSTACK	18 Reserved
-------------------------	----------------

The following DCB operands pertain to specific devices and may be specified only when the DEVD parameter is specified.

KEYLEN=length

for direct access devices, the length in bytes of the key of a physical record, with a maximum value of 255. When a block is read or written, the number of bytes transmitted is the key length plus the record length.

DEN=value

for magnetic tape, the tape recording density in bits per inch:

Value	7-Track Density	9-Track Density
1	556	---
2	800	800(NRZI)
3	---	1600(PE)
4	---	6250(GCR)

NRZI—Non-return-to-zero change to ones recording

PE—phase encoded recording

GCR—group coded recording

If this parameter is omitted, the highest density available on the device is assumed.

TRTCH=value

for 7-track magnetic tape, the tape recording technique:

Value Tape Recording Technique

C Data conversion feature is available.

E Even parity is used. (If omitted, odd parity is assumed.)

T BCDIC to EBCDIC translation is required.

MODE=value

for a card reader or punch, the mode of operation. Either C (column binary mode) or E (EBCDIC code) may be specified.

STACK=value

for a card punch or card reader, the stacker bin to receive cards, either 1 or 2.

PRTSP=value

for a printer, the line spacing, either 0, 1, 2, or 3.

DSORG Parameter of the DCBD Macro

In addition to the operands described in Data Administration: Macro Instruction Reference for the DSORG parameter of the DCBD macro, you may specify the following operands.

DSORG=

XA specifies a DCB with the EXCP interface section (including appendage names)

XE specifies a DCB with the foundation block extension

INITIALIZING DATA CONTROL BLOCKS (OPEN)

The OPEN macro instruction initializes one or more data control blocks so that their associated data sets can be processed. You must issue OPEN for all data control blocks that are to be used by your channel programs. (A dummy data set may not be opened for EXCP.) Some of the procedures performed when OPEN is executed are:

- Reading in the JFCB (job file control block), unless the TYPE=J option of the macro instruction was coded
- Construction of the data extent block (DEB)
- Transfer of information from the JFCB and data set labels to the DCB
- Verification or creation of standard labels
- Tape positioning
- Loading of your appendage routines

The parameters of the OPEN macro instruction are:

<u>[symbol]</u>	OPEN	(<u>dcb address</u> , <u>[options]</u>),...
-----------------	------	--

dcb address—A-type address or (2-12)

the address of the data control block to be initialized. (More than one data control block may be specified.)

option1

the intended method of I/O processing of the data set. You may specify this parameter as either INPUT, RDBACK, OUTPUT, or EXTEND. For magnetic tape, label processing for each of these when OPEN is executed is as follows:

INPUT Header labels are verified.
RDBACK Trailer labels are verified.
OUTPUT Header labels are created.
EXTEND Header labels are created.

If this parameter is omitted, INPUT is assumed.

option2

the volume disposition that is to be provided when volume switching occurs. The operand values and meanings are as follows:

REREAD Reposition the volume to process the data set again.

LEAVE No additional positioning is performed at end-of-volume processing.

DISP Specifies that a tape volume is to be disposed of in the manner implied by the DD statement associated with the data set. Direct access volume positioning and disposition are not affected by this parameter of the OPEN macro instruction. There are several dispositions that can be specified in the DISP parameter of the DD statement:

DISP=PASS, DELETE, KEEP, CATLG, or UNCATLG. Only DISP=PASS has significance at the time an end-of-volume condition is encountered. The end-of-volume condition may result from the issuance of an FEOV macro instruction or may be the result of reaching the end of a volume.

If DISP=PASS was coded in the DD statement, the tape will be spaced forward to the logical end of the data set on the current volume.

If a DISP option other than DISP=PASS is coded on the DD statement, the action taken when an end-of-volume condition occurs depends on (1) how many tape units are allocated to the data set and (2) how many volumes are specified for the data set in the DD statement. This is determined by the UNIT= and VOLUME= operands of the DD statement associated with the data set. If the number of volumes is greater than the number of units allocated, the current volume will be rewound and unloaded. If the number of volumes is less than or equal to the number of units, the current volume is merely rewound.

If you intend to process a multivolume direct data set, you must cause open routines to build a data extent block for each volume and issue mount messages for them. This can be done by reading in the JFCB with a RDJFCB macro instruction and opening each volume of the data set. The following code illustrates the procedure:

```

RDJFCB      DCB1      READS IN THE JFCB
SR          R3,R3    CLEARS REG 3; IT WILL
*           *        HOLD COUNT OF VOLS TO
*           *        BE OPENED
*           IC      R3,JFCBNVOL PUTS # OF VOLS
*           LA      R4,DCB1    IN REG 3
*           LA      R5,1      R4 POINTS TO DCB FOR
*           *        VOL TO BE OPENED
*           *        PUTS SEQUENCE # OF
*           *        FIRST VOL TO BE
LOOP        EQU      *        OPENED IN REG 5
           STH      R5,JFCBVLSQ PUTS SEQ # OF VOL
*           *        TO BE OPENED WHERE
*           *        OPEN RTNS LOOK
*           *        OPENS ONE VOL
           OPEN ((R4),OUTPUT),TYPE=J
*           *        NOTE THAT THE TYPE=J OPTION OF THE MACRO MUST BE USED
           LA      R4,DCB2-DCB1(R4) INCREMENT REG 4 TO
*           *        POINT TO THE DCB FOR
*           *        THE NEXT VOL TO BE
*           *        OPENED
           LA      R5,1(R5)    INCREMENT TO SEQ # OF
*           *        NEXT VOL TO BE OPENED
           BCT     R3,LOOP    LOOP UNTIL ALL VOLS
*           *        OPEN
           .
           .
JFCB        DS      CL176     JFCB READ IN HERE
           ORG     JFCB+70
JFCBVLSQ   DS      H         SEQ # OF VOL TO BE
*           *        OPENED
           ORG     JFCB+117
JFCBNVOL   DS      FL1      # OF VOLS IN DATA SET
           ORG
* MAPPING MACRO IEFJFCBN MAY ALSO BE USED
DCB1 DCB DDNAME=SYSUT1,MACRF=(E),EXLST=EXITS,DSORG=PS
DCB2 DCB DDNAME=SYSUT1,MACRF=(E),EXLST=EXITS,DSORG=PS
DCB3 DCB DDNAME=SYSUT1,MACRF=(E),EXLST=EXITS,DSORG=PS
DCB4 DCB DDNAME=SYSUT1,MACRF=(E),EXLST=EXITS,DSORG=PS
DCB5 DCB DDNAME=SYSUT1,MACRF=(E),EXLST=EXITS,DSORG=PS
* THIS PROCEDURE WORKS FOR 5 VOLS OR LESS; THE JFCB
* EXTENSION, WHICH IDENTIFIES ADDITIONAL VOLS, CAN'T
* BE READ IN
EXITS     DS      OF
           DC      X'87',AL3(JFCB) 87 IDENTIFIES THIS AS
*           *        THE EXIT LIST ENTRY
*           *        THAT SHOWS WHERE JFCB
*           *        WILL BE READ IN

```

Use of the RDJFCB macro instruction and the OPEN macro instruction with the TYPE=J option is explained in detail under "Reading and Modifying a Job File Control Block" on page 117.

EXECUTING A CHANNEL PROGRAM (EXCP)

The EXCP macro instruction requests the initiation of the I/O operations of a channel program. You must issue EXCP whenever you want to execute one of your channel programs. The format of the EXCP macro instruction is:

[symbol]	EXCP	job-address
----------	------	-------------

job-address—A-type address, (2-12), or (1)
the address of the input/output block of the channel program to be executed.

ASSIGNING AN ALTERNATE TRACK AND COPYING DATA FROM THE DEFECTIVE TRACK (ATLAS)

A program that uses the EXCP macro instruction for input and output and that is APF authorized may, during the execution of the program, use the ATLAS macro instruction to obtain an alternate track and to copy a defective track onto the alternate track. With the use of ATLAS, the program can recover from permanent (hard) errors encountered in the execution of the following types of I/O commands:

- Search ID.
- Write. (The error condition must be confirmed during the execution of the channel program by a CCW that checks the data written.)
- Read count. Errors in the CCHHR part of the count area can be recovered from, unless the record is the home address or record zero. Errors in the KDD part of the count area cannot be recovered from, unless the user has identified the defective record.

Note: ATLAS may be used for all direct access devices with the exception of MSS volumes (3330V).

Your DCB must include the DCBRECFCM field, and the field must show whether the data set is in the track overflow format. If it is, recovery from errors in last records on tracks depends on your identifying the track overflow record segments.

Recovery takes the form of obtaining a good alternate track and copying the defective track onto the good alternate one. Unless a reexecution of the channel program by ATLAS can correct the defect, the user should examine, and if necessary replace, defective records in a subsequent job if the data set is to be processed again.

The format is:

<u>[symbol]</u>	ATLAS	PARMADR={ <u>address</u> } [,CHANPRG={R def. NR}] [,CNTPTR={P F}] [,WRITR={YES NO}]
-----------------	-------	--

PARMADR

Address of a parameter address list of the following format:

0	Address of IOB for the channel program that encountered the error
4	Address of count area field

The count area field contains the CCHHRKDD of a defective record or the CCHH of a track that is to be copied.

address—A-type address, (2-12), or (1)

CHANPRG={R|NR}

specifies whether the channel program that encountered the error can be executed again.

R Channel program may be executed again by ATLAS. Before permitting reexecution of the channel program by ATLAS, you must reset the error indications of the previous execution fields in the DCBIFLGS. (See the example of the use of ATLAS below.)

NR Channel program may not be executed again.

If this parameter is omitted, R is assumed.

CNTPTR

specifies whether the count area field contains a full count area (CCHHRKDD) or a partial count area (CCHH).

P Part of the count area (the CCHH address of the track to be copied).

F Full count area (CCHHRKDD count of the record that was found defective).

If this parameter is omitted, P is assumed.

WRITS

track overflow segment identification.

If your data set is in the track overflow format, this identification determines recovery from errors in last records on tracks.

YES If this is the last record on the track, it is a segment other than the last of a track overflow record.

NO If this is the last record on the track, it is the last or only segment of a track overflow record.

If this parameter is omitted, it is assumed that it cannot be established whether a last record is a segment of an overflow record.

Using ATLAS

If a channel program encounters a unit check condition (shown in the CSW) in its execution, the EXCP Processor program will place the sense bytes in the IOB. ATLAS can be used to recover from sense conditions shown by the following bit settings:

IOBSENS0 X'08' Data check

IOBSENS1 X'80' Permanent

Also, before using ATLAS, you must reset error indications as follows:

NI DCBIFLGS,X'3F' Reset the DCBIFLGS error indications.

The ATLAS program will attempt to find a good alternate track and will attempt to copy the defective track onto the good track, including all error conditions in either key or data areas. The error conditions may be rectified by reexecuting the channel program or through the use of the IEHATLAS utility program in a subsequent step.

Example: The following illustrates the use of the ATLAS macro instruction.

	EXCP	MYIOB	
	WAIT	ECB=MYECB	
	TM	MYECB,X'7F'	TEST FOR I/O ERROR
*	BO	NEXT	NO, SUCCESSFUL, GO TO
			ANOTHER ROUTINE
	TM	IOBCSW+3,X'02'	UNIT CHECK
*	BZ	OTHER	NO, DO OTHER ERROR
			PROCESSING
	TM	IOBSENS0,X'08'	DATA CHECK
	BNO	OTHER	NO, CAN'T HANDLE
	TM	IOBSENS1,X'80'	PERMANENT
	BNO	OTHER	NO, CAN'T HANDLE
*	NI	DCBIFLGS,X'3F'	RESET ERROR
			INDICATORS
	ATLAS	PARMADR=THREE,CHANPRG=R	

Operation of the ATLAS Program

The ATLAS program (SVC 86):

- Establishes the availability and address of the next alternate track from the format-4 DSCB of the VTOC.
- Brings all count fields from the defective track into storage to establish the description of the track.
- Initializes the alternate track. (Writes the home address and record zero.)
- Brings the key and data areas of each record into storage, one at a time, and combines them with their new count area to write the complete record onto the alternate track.
- When the copying is finished, chains the alternate to the defective track and updates the VTOC.

Control is returned to your program at the next executable instruction following the ATLAS macro instruction.

Return Codes from the ATLAS Program

The success of the ATLAS macro instruction can be determined by examining the contents of register 15, which will contain one of the return codes described below. If register 15 contains decimal 0, 36, 40, or 44, the contents of register 0 may be significant.

Code	Meaning
0(X'00')	<p>Successful completion. Key and data areas have been copied from the defective track onto a good alternate one. The only error encountered was in the record identified by the user's CCHHRKDD value.</p> <p>If the channel program is reexecutable, it has been successfully reexecuted.</p>
4(X'04')	This device type does not have alternate tracks that can be assigned by programming.
8(X'08')	All alternate tracks for the device have been assigned.
12(X'0C')	A request for storage (GETMAIN macro instruction) could not be satisfied.
16(X'10')	All attempts to initialize and transfer data to an alternate track failed. The number of attempts made is equal to 10% of the assigned alternates for the device.
20(X'14')	The type of error shown by the sense byte cannot be handled through the use of the ATLAS macro instruction. The condition is other than a data check (in the count or data areas) or a missing address marker.
24(X'18')	The format-4 DSCB of the VTOC cannot be read; therefore alternate track information is not available to ATLAS.
28(X'1C')	The record specified by the user was the format-4 DSCB, and it could not be read.
32(X'20')	An error found in count area of last record on the track cannot be handled because last-record-on-track identification is not supplied.
36(X'24')	<p>An error was encountered when reading or writing the home address record or record zero. No error recovery has taken place.</p> <p>If register 0 contains X'01 00 00 00', the defect is in record zero.</p>
40(X'28')	<p>Successful completion. Key and data areas have been copied from the defective track onto a good alternate one. However, the alternate track may have records with defective key or data areas. Register 0 identifies the first three found defective as follows:</p> <p style="padding-left: 40px;">n R R R</p> <p>n—The number of record numbers that follow (0, 1, 2, or 3).</p> <p>R—The hexadecimal number of the record found defective but copied anyway.</p> <p>If the channel program is reexecutable, it has been successfully reexecuted.</p>

Code	Meaning
44(X'2C')	<p>Errors encountered and no alternate track has been assigned. The return parameter register (register 0) will contain the R of a maximum of three error records.</p> <p>Error conditions that return this code are:</p> <ol style="list-style-type: none"> 1. ATLAS received an error indication for a record with a data length in the count field of zero. Recovery was not possible because a distinction cannot be made between an EOF record and an invalid data length. 2. An error occurred while reading the count field of a record, and the KDD (key length-data length) was found to be defective. 3. More than three records on the specified track contained errors in their count fields.
48(30)	No errors found on the track, no alternate assigned. ATLAS will not assign an alternate unless a track has at least one defective record.
52(34)	I/O error in reexecuting user's channel program. A good alternate is chained to the defective track, and data has been transferred. The user's control blocks will give indication of the error condition causing failure in reexecution of the channel program.
56(38)	The DCB reflects a track overflow data set, but the UCB device type shows that the device does not support track overflow.
60(X'3C')	The CCHH of the user-specified count area is not within the extents of the data set.
64(X'40')	The device is an MSS virtual device, which is not supported.

END OF VOLUME (EOV)

The EOV macro instruction identifies end-of-volume and end-of-data-set conditions. For an end-of-volume condition, EOV causes switching of volumes and verification or creation of standard labels. For an end-of-data-set condition, EOV causes your end-of-data set routine to be entered. Before processing trailer labels on a tape input data set, you must decrement the DCBBLKCT field. You issue EOV if switching of magnetic tape or direct access volumes is necessary, or if secondary allocation is to be performed for a direct access data set opened for output.

For magnetic tape, you must issue EOV when either a tapemark is read or a reflective spot is written over. In these cases, bit settings in the 1-byte DCBOFLGS field of the data control block determine the action to be taken when EOV is executed. Before issuing EOV for magnetic tape, you must make sure that appropriate bits are set in DCBOFLGS. Bit positions 2, 3, 6, and 7 of DCBOFLGS are used only by the system; you are concerned with bit positions 0, 1, 4, and 5. The use of these DCBOFLGS bit positions is as follows:

- Bit 0 set to 1 indicates that a write command was executed and that a tapemark is to be written.
- Bit 1 indicates that a backward read was the last I/O operation.
- Bit 4 indicates that data sets of unlike attributes are to be concatenated.
- Bit 5 indicates that a tapemark has been read.

If bits 0 and 5 of DCBOFLGS are both off when EOVS is executed, the tape is spaced past a tapemark, and standard labels, if present, are verified on both the old and new volumes. The direction of spacing depends on bit 1. If bit 1 is off, the tape is spaced forward; if bit 1 is on, the tape is backspaced.

If bit 0 is on, but bit 5 is off, when EOVS is executed, a tapemark is written immediately following the last data record of the data set. Standard labels, if specified, are created on the old and the new volume.

After issuing EOVS for sequentially organized output data sets on direct access volumes, you can determine whether additional space was obtained on the same or a different volume. You do this by examining the data extent block (DEB) and the unit control block (UCB). If neither the address of the UCB, as shown in the DEB, nor the volume serial number, as shown in the UCB, has changed, additional space was obtained on the same volume. Otherwise, space was obtained on a different volume.

The only parameter of the EOVS macro instruction is:

[symbol]	EOVS	dcb address
----------	------	-------------

dcb address—A-type address, (2-12), or (1)
the address of the data control block that is opened for the data set. If this parameter is specified as (1), register 1 must contain this address.

Note: To learn how the system disposes of a tape volume when an EOVS macro is issued, see the description of the DISP parameter under "Initializing Data Control Blocks (OPEN)" on page 58.

RESTORING DATA CONTROL BLOCKS (CLOSE)

The CLOSE macro instruction restores one or more data control blocks so that processing of their associated data sets can be terminated. You must issue CLOSE for all data control blocks that were used by your channel programs. Some of the procedures performed when CLOSE is executed are:

- Release of data extent block (DEB)
- Removal of information transferred to data control block fields when OPEN was executed
- Verification or creation of standard labels
- Volume disposition
- Release of programmer-written appendage routines

When CLOSE is issued for data sets on magnetic tape volumes, labels are processed according to bit settings in the DCBOFLGS field of the data control block. Before issuing CLOSE for magnetic tape, you must set the appropriate bits in DCBOFLGS.

The DCBOFLGS bit positions that you are concerned with are listed in the EOVS macro instruction description.

For information about the forms of the CLOSE macro and their parameters, see Data Administration: Macro Instruction Reference.

CONTROL BLOCK FIELDS

The fields of the input/output block, event control block, and data extent block are illustrated and explained here; the data control block fields are described with the parameters of the DCB macro instruction under "EXCP Requirements" on page 39.

INPUT/OUTPUT BLOCK (IOB) FIELDS

The input/output block (IOB) is not automatically constructed by a macro instruction; it must be defined as a series of constants and must be on a fullword boundary. For unit-record and tape devices, the IOB is 32 bytes in length. For direct access, teleprocessing, and graphic devices, 8 additional bytes must be provided. You may want to use the system mapping macro IEZIOB, which expands into a DSECT, to help in constructing an IOB.

In Figure 14 the diagonally-ruled areas indicate fields in which you must specify information. The other fields are used by the system and must be defined as all zeros. You may not place information into these fields, but you may examine them.

IOBFLAG1 (1 byte)

You must set bit positions 0, 1, and 6. One-bits in positions 0 and 1 indicate data chaining and command chaining, respectively. (If both data chaining and command chaining are specified, the system does not use error recovery routines except for the direct access devices.) A one-bit in position 6 indicates that the channel program is not a "related" request; that is, the channel program is not related to any other channel program. If you intend to issue an EXCP macro with a BSAM, QSAM, or BPAM data control block, you may want to turn on bit 7 to prevent access-method appendages from processing the I/O request.

IOBFLAG2 (1 byte)

If you set bit 6 in the IOBFLAG1 field to zero, bits 2 and 3 in this field must then be set to:

- 00, if any channel program or appendage associated with a related request might modify this IOB or channel program.
- 01, if the conditions requiring a 00 setting don't apply, but the CHE or ABE appendage might retry this channel program if it completes normally or with the unit-exception or wrong-length-record bits on in the CSW.
- 10 in all other cases.

The three combinations of bits 2 and 3 represent the three kinds of related requests, known as type 1 (00), type 2 (01), and type 3 (10). The type you use determines how much the EXCP Processor can overlap the processing of related requests. Type 3 allows the greatest overlap, normally making it possible to quickly reuse a device after a channel-end interruption. (Related requests that were executed on a pre-MVS system are executed as type-1 requests if not modified.)

IOBSENS0 and IOBSENS1 (2 bytes)

are placed into the input/output block by the EXCP Processor when a unit check occurs. On occasion, the system is unable to obtain any sense bytes because of unit

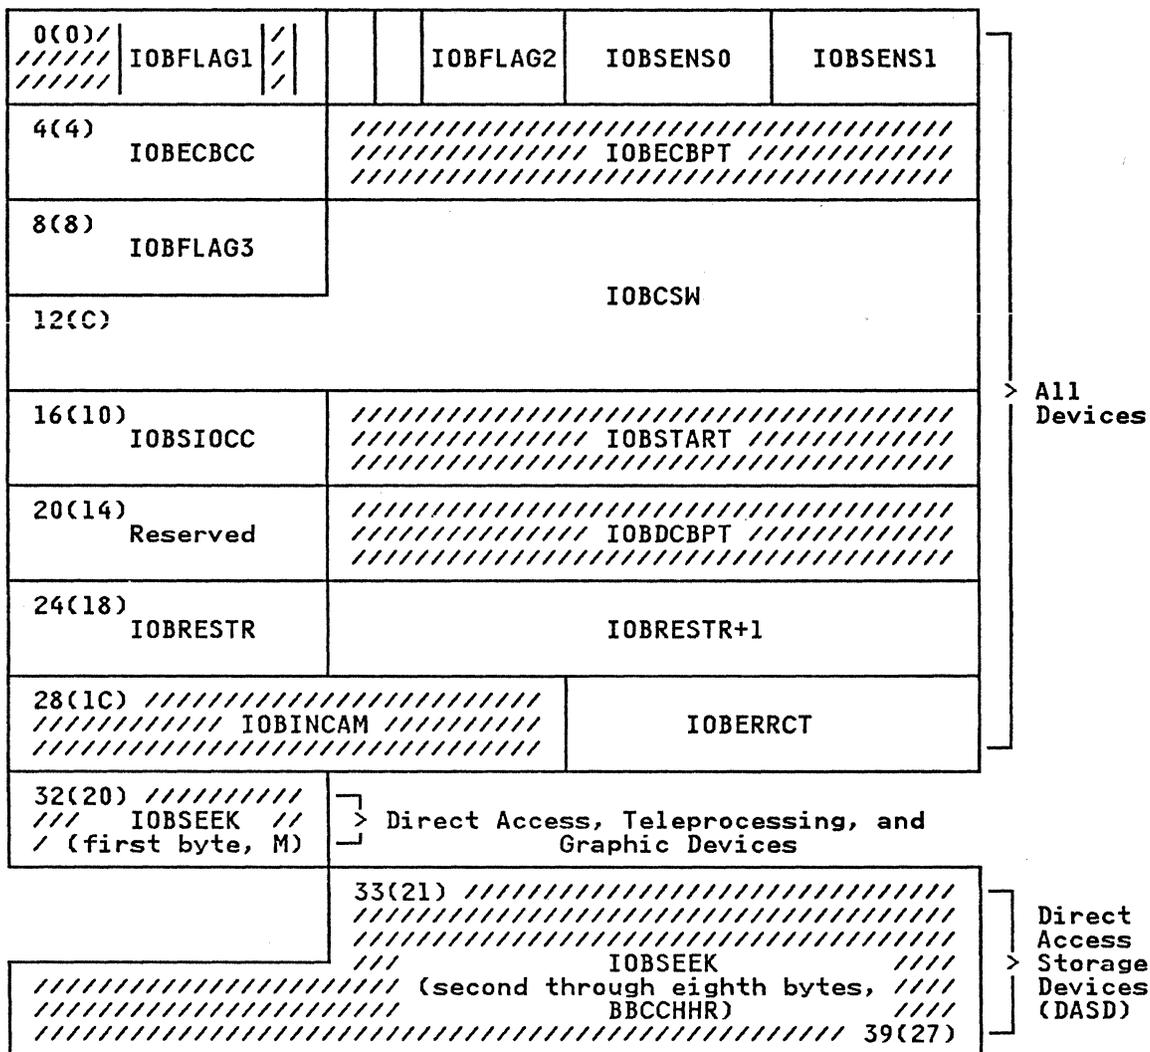


Figure 14. Input/Output Block (IOB) Format

checks when sense commands are issued. In this case, the system simulates sense bytes by moving X'10FE' to IOBSENS0 and IOBSENS1.

IOBECBCC (1 byte)
the first byte of the completion code for the channel program. The system places this code in the high-order byte of the event control block when the channel program is posted complete. The completion codes and their meanings are listed under "Event Control Block (ECB) Fields" on page 69.

IOBECBPT (3 bytes)
the address of the 4-byte event control block you have provided.

IOBFLAG3 (1 byte)
is used only by the system.

IOBCSW (7 bytes)
the low-order seven bytes of the channel status word that are placed into this field each time a channel-end or PCI interruption occurs.

IOBSIOCC (1 byte)
in bits 0 and 1, the instruction-length code; in bits 2 and 3, the start subchannel (SSCH) condition code for the instruction the system issues to start the channel program; and, in bits 4 through 7, the program mask.

IOBSTART (3 bytes)
the starting address of the channel program to be executed.

Reserved (1 byte)
used only by the system.

IOBDCBPT (3 bytes)
the address of the data control block of the data set to be read or written by the channel program.

IOBRESTR (1 byte)
used by the system for volume repositioning in error recovery procedures.

IOBRESTR+1 (3 bytes)
if a related channel program is permanently in error, used by the system to chain together IOBs that represent dependent channel programs. To learn more about the conditions under which the chain is built, see "Interruption Handling and Error Recovery Procedures" on page 42.

IOBINCAM (2 bytes)
for magnetic tape, the amount by which the block count (DCBBLKCT) field in the device-dependent portion of the data control block is to be incremented. You may alter these bytes at any time. For forward operations, these bytes should contain a binary positive integer (usually +1); for backward operations, they should contain a binary negative integer. When these bytes are not used, all zeros must be specified.

Reserved (2 bytes)
used only by the system.

IOBSEEK (first byte, M)
for direct access devices, the extent entry in the data extent block that is associated with the channel program (0 indicates the first entry; 1 indicates the second, and so forth). For teleprocessing and graphic devices, it contains the UCB index.

IOBSEEK (last 7 bytes, BBCCHHR)
for direct access devices, the seek address for your channel program.

EVENT CONTROL BLOCK (ECB) FIELDS

You must define an event control block (ECB) as a 4-byte area on a fullword boundary. When the channel program has been completed, the input/output supervisor places a completion code containing status information into the ECB (Figure 15 on page 70). Before examining this information, you must test for the setting of the "complete bit." If the complete bit is not on, and your problem program cannot perform other useful operations, you should issue a WAIT macro instruction that specifies the event control block. Under no circumstances should you construct a program loop that tests for the complete bit.

eliminates the translation of channel programs by the EXCP processor. The program issuing the EXCPVR must remain in authorized state until the completion of the channel programs.

Problem programs are authorized to use the EXCPVR macro instruction under the authorized program facility (APF). A description of how to authorize a program can be found in Supervisor Services and Macro Instructions.

<code>[symbol]</code>	EXCPVR	<code>job-address</code>
-----------------------	--------	--------------------------

job-address—A-type address, (2-12), or (1)
the address of the input/output block of the channel program to be executed.

To use EXCPVR, you must do all the things you would do to execute an EXCP request; in addition you must:

1. Code PGFX=YES in the DCB associated with the EXCPVR requests and provide a page-fix (PGFX) appendage by specifying SIOA=symbol in the DCB.
2. Fix the data area that contains your channel program, the data areas that are referred to by your channel program, your PCI appendage (if your program can generate program-controlled interrupts), and any area referred to by your PCI appendage. To cause EXCP to fix these data areas, you build a list that contains the addresses of these virtual areas. You should build the list in your PGFX appendage.
3. Determine whether the data areas in virtual storage specified in the address fields of your CCWs cross page boundaries. If they do, you must build an indirect data address list (IDAL) and put the address of the IDAL in the affected CCW.
4. Translate the addresses in your CCWs from virtual to real addresses.

All other areas related to the EXCP/EXCPVR I/O operation (that is, CCWs, IDAWs, IOBs, DEBs, DCBs, appendages, and so forth) must remain 24-bit addressable. Note, however, that the EXCP processor will allow both 24-bit and 31-bit virtual I/O buffers to be fixed above 16 megabytes real.

Items 3 and 4 must be done in your start-I/O (SIO) appendage. A description of the SIO appendage is presented under "Appendages" on page 43.

BUILDING THE LIST OF DATA AREAS TO BE FIXED

The EXCP processor expects programs using the EXCPVR macro instruction to pass a list of data areas to be fixed. This list is to be built in the PGFX appendage, as described below.

The data areas you must fix in real storage (if not already fixed in real storage) are:

1. The channel program. If the channel program is already in a fixed subpool, it does not have to be fixed.
2. The data areas to which your channel program will be writing and from which your channel program will be reading. If the data areas are already in a fixed subpool, they do not have to be fixed.
3. The PCI appendage, if used, and any areas referred to in the PCI appendage.
4. Any system or user control blocks (and the DEB).

You need not fix areas that have already been fixed, such as the modules that reside in the fixed link pack area (LPA).

EXCPVR users can specify 31-bit real data areas by creating CCWs through the use of IDAWs.

PAGE FIX (PGFX) AND START-I/O (SIO) APPENDAGE

This appendage comprises two essentially independent appendages. The complete appendage can be viewed as a re-entrantable subroutine having two entry points, one for the SIO appendage and one for the PGFX appendage.

The SIO entry point is located at offset 0 in the subroutine; any other location in the appendage may be branched to from this entry point. The entry point of the PGFX appendage is at offset +4 in the SIO subroutine, which is set in register 15 as the entry point of the PGFX appendage.

Page Fix (PGFX) Appendage: The purpose of this appendage is to list all the areas that must be fixed to prevent paging exceptions during the execution of the current I/O request. This appendage may be entered more than once. However, each time it is entered, it must create the same list of areas to be fixed. The appendage may use the 16-word save area pointed to by register 13. Registers 10, 11, and 13 may be used as work registers.

Page-Fix List Processing

Each page-fix entry placed in the list by the appendage must have the following doubleword format:

0 1	31 32	33	63
0	Starting virtual address of area to be fixed	0	Ending virtual address of area to be fixed + 1

On return from your PGFX appendage to the EXCP processor (via the return address provided in register 14), register 10 must point to the first page-fix entry and register 11 must contain the number of page-fix entries in the work area. The EXCP processor then fixes the pages corresponding to the areas listed by the PGFX appendage. The pages remain fixed until the associated I/O request terminates.

If either the channel end appendage or the abnormal end appendage returns via the return address in register 14 plus 8, the PGFX appendage is not normally reentered. Instead, the SIO appendage is entered, and the page-fix list built by the PGFX appendage is still active. However, the PGFX appendage is entered after either the channel end appendage or the abnormal end appendage returns via the return address in register 14 plus 8 when a PURGE macro has been issued (for instance, when a storage swap has occurred). In this case, when I/O is restored, the PGFX appendage is entered.

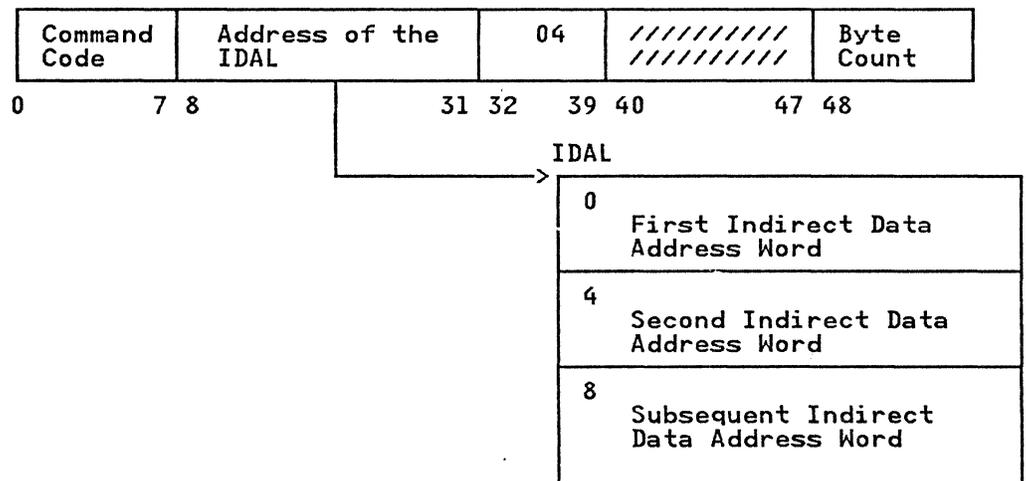
Note: The page-fix list must be in page-fixed storage.

SIO APPENDAGE: If you are using EXCPVR to execute your channel program, you must translate the virtual addresses in the operands of your channel program to real addresses. This should be done in your SIO appendage. If indirect data addressing is required, the SIO appendage should also build the indirect data address lists (IDALs) and turn on the IDA indicators in the associated CCWs.

Translating Virtual Addresses and Building the IDAL: You must convert the virtual addresses in the channel program to real addresses. You must also check the areas whose addresses appear in bits 8 through 31 of your CCWs to determine whether the data areas cross 2K-byte boundaries. If they do, you must provide an entry in the IDAL for each 2K-byte boundary crossed. The channel subsystem uses the IDAL to identify the address where it will continue reading or writing when a 2K-byte boundary is crossed during a read or write operation. The IDAL must contain real addresses when it is processed by the channel.

In MVS/XA, the LRA instruction returns a 31-bit real address regardless of the addressing mode. You must be careful when you construct the IDAW to ensure that the real storage obtained by GETMAIN (or branch entry) is below 16 megabytes. Do your page fixing before you issue the LRA. (See Supervisor Services and Macro Instructions or System Macros and Facilities for information on how to obtain real storage below 16 megabytes real.)

CCW



Notes:

1. You must put one entry in the IDAL for each 2K-byte page boundary your data area crosses.
2. If the CCW has an IDAL address rather than a data address, bit 37 must be set to signal this to the channel.
3. The maximum number of entries needed in the IDAL is determined from the count in the CCW as follows:

Number of IDAL entries = ((CCW byte-count - 1) / 2048) + 1.
(Round up division to next highest integer if remainder is not zero.)

The number of IDAL entries required ultimately depends on the number of 2K-byte boundaries crossed by the data. For example, if your data is 800 bytes long and does not cross a 2K-byte page boundary, no IDAL entries are required. If your data crosses a 4K-byte page boundary, then two IDAL entries are required. If your data is 5000 bytes long, at least two IDAL entries are required. If your data crosses two 4K-byte page boundaries, four IDAL entries are required.

The first indirect address is the real address of the first byte of the data area. The second and subsequent indirect addresses are the real addresses of the second and subsequent 2K-byte boundaries of the data area.

For example, if the data area real address is X'707FF' and the byte count is X'1802', the IDAL would contain the following real addresses (assuming the real addresses are contiguous, which may not always be the case):

707FF
70800
71000

If the data area real address is X'707FF' and the byte count is X'800', the IDAL would contain the following addresses:

707FF
70800

CHAPTER 3. READING FROM AND WRITING TO DIRECT ACCESS DEVICES (XDAP)

Execute direct access program (XDAP) is a macro instruction that you may use to read, verify, or update a block on a direct access volume. This chapter explains what the XDAP macro instruction does and how you can use it. The control block generated when XDAP is issued and the macro instructions used with XDAP are also discussed.

Because most of the specifications for XDAP are similar to those for the execute channel program (EXCP) macro instruction, you should be familiar with the "Executing Your Own Channel Programs (EXCP)" chapter of this publication, and with the information contained in Data Administration Guide that provides how-to information for using the access method routines of the system control program.

If you are not using the standard IBM data access methods, you can, by issuing XDAP, generate the control information and channel program necessary for reading or updating the records of a data set. (However, XDAP cannot be used to read, verify, or update a SYSIN or SYSOUT data set.)

You cannot use XDAP to add blocks to a data set, but you can use it to change the keys of existing blocks. Any block configuration and any data set organization can be read or updated.

Although the use of XDAP requires less storage than do the standard access methods, it does not provide many of the control program services that are included in the access methods. For example, when XDAP is issued, the system does not block or unblock records and does not verify block length.

To issue XDAP, you must provide the actual track address of the track containing the block to be processed. You must also provide either the block identification or the key of the block, and specify which of these is to be used to locate the block. If a block is located by identification, both the key and data portions of the block may be read or updated. If a block is located by key, only the data portion can be processed.

For additional control over I/O operations, you may write appendages that must be entered into the LPA library. Descriptions of these routines and their coding specifications are included under Chapter 2, "Executing Your Own Channel Programs (EXCP)" on page 36.

XDAP REQUIREMENTS

When using the XDAP macro instruction, you must, somewhere in your program, code a DCB macro instruction that produces a data control block (DCB) for the data set to be read or updated. You must also code an OPEN macro instruction that initializes the data control block and produces a data extent block (DEB). The OPEN macro instruction must be executed before any XDAP macro instructions are executed.

When the XDAP macro instruction is assembled, a control block and executable code are generated. This control block may be logically divided into three sections:

- An event control block (ECB) that is supplied with a completion code each time the direct access channel program is terminated.
- An input/output block (IOB) that contains information about the direct access channel program.
- A direct access channel program that consists of three or four channel command words (CCWs). The type of channel program generated depends on specifications in the parameters of the XDAP macro instruction. When executed, it locates a block by either its actual address or its key and reads, updates, or verifies the block.

When the channel program has terminated, a completion code is placed into the event control block. After issuing XDAP, you should therefore issue a WAIT macro instruction, specifying the address of the event control block, to regain control when the direct access program has terminated. If volume switching is necessary, you must issue an EOVS macro instruction. When processing of the data set has been completed, you must issue a CLOSE macro instruction to restore the data control block.

MACRO SPECIFICATIONS FOR USE WITH XDAP

When you are using the XDAP macro instruction, you must also code DCB, OPEN, CLOSE, WAIT, and, in some cases, the EOVS macro instructions. The parameters of the XDAP macro instruction are listed and described here. For the other required macro instructions, special requirements or options are explained, but you should see "Macro Specifications for Use with EXCP" on page 51 for listings of their parameters.

DEFINING A DATA CONTROL BLOCK (DCB)

You must issue a DCB macro instruction for each data set to be read, updated, or verified by the direct access channel program. To learn which macro instruction parameters to code, see "Defining Data Control Blocks for EXCP (DCB)" on page 51.

INITIALIZING A DATA CONTROL BLOCK (OPEN)

The OPEN macro instruction initializes one or more data control blocks so that their associated data sets can be processed. You must issue OPEN for all data control blocks that are to be used by the direct access program. Some of the procedures performed when OPEN is executed are:

- Construction of data extent block (DEB)
- Transfer of information from DD statements and data set labels to the data control block
- Verification or creation of standard labels
- Loading of programmer-written appendage routines

The two parameters of the OPEN macro instruction are the address(es) of the data control block(s) to be initialized and the intended method of I/O processing of the data set. The method of processing may be specified as INPUT, OUTPUT, EXTEND; however, if nothing is specified, INPUT is assumed.

| EXECUTING DIRECT ACCESS PROGRAMS (XDAP)

The XDAP macro instruction produces the XDAP control block (that is, the ECB, IOB, and channel program) and executes the direct access channel program. The format of the XDAP macro instruction is:

[symbol]	XDAP	<u>ecb-symbol</u> , <u>type</u> , <u>dcb-addr</u> , <u>area-addr</u> , <u>length-value</u> , <u>[[key-addr, keylength-value]]</u> , <u>blkref-addr</u> , <u>[sector-addr]</u> , <u>[MF={E L}]</u>
----------	------	---

ecb-symbol—symbol or (2-12)

the symbolic name to be assigned to the XDAP event control block. Registers can be used only with MF=E.

type—{RI|RK|WI|WK|VI|VK}

the type of I/O operation intended for the data set and the method by which blocks of the data set are to be located. One of the combinations shown must be coded in this field.

The codes and their meanings are:

- R Read a block.
- W Update a block.
- V Verify that the device is able to read the contents of a block, but do not transfer data.
- I Locate a block by identification. (The key portion, if present, and the data portion of the block are read, updated, or verified.)
- K Locate a block by key. (Only the data portion of the block is read, updated, or verified.) If you code this value, you must code the 'key-addr, keylength-value' operands.

dcb-addr—A-type address or (2-12)

the address of the data control block for the data set. If this data control block is also being used by a sequential access method (BSAM, BPAM, QSAM), you must reassemble the XDAP macro instruction. Otherwise, sequential access method appendages will be called at the conclusion of the XDAP channel program.

area-addr—A-type address or (2-12)

the address of an input or output area for a block of the data set.

length-value—absexp or (2-12)

the number of bytes to be transferred to or from the input or output area. If blocks are to be located by identification and the data set contains keys, the value must include the length of the key. The maximum number of bytes transferred is 32767.

key-addr—RX-type address or (2-12)

when blocks are to be located by key, the address of a virtual storage field that contains the key of the block to be read, updated, or verified.

keylength-value—absexp or (2-12)

when blocks are to be located by key, the length of the key. The maximum length is 255 bytes.

blkref-addr—RX-type address or (2-12)

the address of a field in virtual storage containing the actual track address of the track containing the block to be located. The actual address of a block is in the form MBBCCHHR, where M indicates which extent entry in the data extent block is associated with the direct access program; BB is not used, but must be zero; CC indicates the cylinder address; HH indicates the actual track address; and R indicates the block identification. R is not used when blocks are to be located by key. (For more detailed information, see "Converting a Relative Track Address to an Actual Track Address" on page 80.)

sector-addr—RX-type address or (2-12)

the address of a 1-byte field containing a sector value. The sector-address parameter is used for rotational position sensing (RPS) devices only. The parameter is optional, but its use will improve channel performance. When the parameter is coded, a set-sector CCW (using the sector value indicated by the data address field) precedes the search-ID-equal command in the channel program. The sector-address parameter is ignored if the type parameter is coded as RK, WK, or VK. If a sector address is specified in the execute form of the macro, then a sector address, not necessarily the same, must be specified in the list form. The sector address in the executable form will be used.

Note: No validity check is made on either the address or the sector value when the XDAP macro is issued. However, a unit check/command reject interruption will occur during channel-program execution if the sector value is invalid for the device or if the sector-addr operand is used when accessing a device without RPS. (For more detailed information, see "Obtaining Sector Number of a Block on a Device with the RPS Feature" on page 82.)

MF=

you may use the L-form of the XDAP macro instruction for a macro expansion consisting of only a parameter list, or the E-form for a macro expansion consisting of only executable instructions.

MF=E

The first operand (ecb-symbol) is required and may be coded as a symbol or supplied in registers 2 through 12. The type, dcb-addr, area-addr, and length-value operands may be supplied in either the L- or E-form. The blkref-addr operand may be supplied in the E-form or moved into the IOBSEEK field of the IOB by you. The sector-addr is optional; it may be coded either in both the L- and E-form or in neither.

MF=L

The first two operands (ecb-symbol and type) are required and must be coded as symbols. If you choose to code length-value or keylength-value, they must be absolute expressions. Other operands, if coded, must be A-type addresses. (blkref-addr is ignored if coded.)

Note: The XDAP macro builds a channel program containing A-type addresses. These addresses refer to storage within the L-form of the macro. If you copy the L-form of the macro to a workarea so that the program may be reentrant, the E-form of the XDAP macro does not update the A-type addresses. This results in an invalid channel program.

The dcb-addr, area-addr, blkref-addr, and sector-value operands may be coded as RX-type addresses or supplied in registers 2 through 12. The length-value and keylength-value operands can be specified as absolute expressions or decimal integers or supplied in registers 2 through 12.

END OF VOLUME (EOV)

The EOV macro instruction identifies end-of-volume and end-of-data-set conditions. For an end-of-volume condition, EOV causes switching of volumes and verification or creation of standard labels. For an end-of-data-set condition, EOV causes your end-of-data-set routine to be entered. When using XDAP, you issue EOV if switching of direct access volumes is necessary or if secondary allocation is to be performed for a direct access data set opened for output.

The only parameter of the EOV macro instruction is the address of the data control block of the data set.

RESTORING A DATA CONTROL BLOCK (CLOSE)

The CLOSE macro instruction restores one or more data control blocks so that processing of their associated data sets can be terminated. You must issue CLOSE for all data sets that were used by the direct access channel program. Some of the procedures performed when CLOSE is executed are:

- Release of data extent block (DEB)
- Removal of information transferred to data control block fields when OPEN was executed
- Verification or creation of standard labels
- Release of programmer-written appendage routines

The CLOSE macro instruction must identify the address of at least one data control block to be restored, and may specify other options. To learn what these options are and how they are specified, see Data Administration: Macro Instruction Reference.

CONTROL BLOCKS USED WITH XDAP

The three control blocks generated during execution of the XDAP macro instruction are described here.

EVENT CONTROL BLOCK (ECB)

The event control block (ECB) begins on a fullword boundary and occupies the first 4 bytes of the XDAP control block. Each time the direct access channel program terminates, the I/O supervisor places a completion code containing status information into the event control block (Figure 16 on page 80). Before examining this information, you must wait for the completion of the channel program by issuing a WAIT macro instruction that specifies the address of the event control block.

INPUT/OUTPUT BLOCK (IOB)

The input/output block (IOB) is 40 bytes in length and immediately follows the event control block. "Control Block Fields" on page 67 contains a diagram of the input/output block (Figure 16 on page 80). You may want to examine the IOBSENS0, IOBSENS1, and IOBCSW fields if the ECB is posted with X'41'.

Type of I/O Operation	CCW	Command Code
Read by identification	1	Search ID Equal
	2	Transfer in Channel
	3	Read Key and Data
Verify by identification ¹	1	Search Key Equal
	2	Transfer in Channel
	3	Read Data
Write by identification	1	Search ID Equal
	2	Transfer in Channel
	3	Write Key and Data
Write by key	1	Search Key Equal
	2	Transfer in Channel
	3	Write Data

¹ For verifying operations, the third CCW is flagged to suppress the transfer of information to virtual storage.

Figure 17. The XDAP Channel Programs

The conversion routine does all its work in general registers. You must load registers 0, 1, 2, 14, and 15 with input to the routine. Register usage is as follows:

Register Use

- 0 Must be loaded with a 4-byte value of the form TTRN, where TT is the track number relative to the beginning of the data set, R is the the block identification on that track, and N is the concatenation number of a BPAM data set. (0 indicates the first data set in the concatenation, a nonconcatenated BPAM data set, or a non-BPAM data set.)
- 1 Must be loaded with the address of the data extent block (DEB) of the data set.
- 2 Must be loaded with the address of an 8-byte area that is to receive the actual address of the block to be processed. The converted address is of the form MBBCCHHR, where M indicates which extent entry in the data extent block is associated with the direct access program (0 indicates the first extent, 1 indicates the second, and so forth); BB is two bytes of zeros; CC is the cylinder address; HH is the actual track address; and R is the block number.
- 3-8 Are not used by the conversion routine.
- 9-13 Are used by the conversion routine and are not restored.
- 14 Must be loaded with the address where control is to be returned after execution of the conversion routine.
- 15 Is used by the conversion routine as a base register and must be loaded with the address where the conversion routine is to receive control.

RETURN CODES FROM THE CONVERSION ROUTINE

When control is returned to your program, register 15 will contain one of the following return codes:

Code	Meaning
0(X'00')	Successful conversion.
4(X'04')	The relative block address converts to an actual track address outside the extents defined in the DEB.

CONVERTING AN ACTUAL TRACK ADDRESS TO A RELATIVE TRACK ADDRESS

To get the relative track address when you know the actual track address, you can use the conversion routine labeled IECPRLTV. The address of the entry point (CVTPRLTV) is in the communication vector table (CVT). The address of the CVT is in location 16.

The conversion routine does all its work in general registers. You must load registers 1, 2, 14, and 15 with input to the routine. Register usage is as follows:

Register	Use
0	Will be loaded with the resulting TTR0 to be passed back to the caller.
1	Must be loaded with the address of the data extent block (DEB) of the data set.
2	Must be loaded with the address of an 8-byte area containing the actual address to be converted to a TTR. The actual address is of the form MBBCCHHR.
3-8	Are not used by the conversion routine.
9-13	Are used by the conversion routine and are not restored.
14	Must be loaded with the address where control is to be returned after execution of the conversion routine.
15	Is used by the conversion routine as a base register and must be loaded with the address where the conversion routine is to receive control.

| OBTAINING SECTOR NUMBER OF A BLOCK ON A DEVICE WITH THE RPS FEATURE

To obtain the performance improvement given by rotational position sensing, you should specify the sector-addr parameter in the XDAP macro. For programs that can be used with both RPS and non-RPS devices, the UCBRPS bit (bit 3 at an offset of 17 bytes into the UCB) should be tested to determine whether the device has rotational position sensing. If the UCBRPS bit is off, a channel program with a "set sector" command must not be issued to the device.

The sector-addr parameter on the XDAP macro specifies the address of a 1-byte field in your region. You must store the sector number of the block to be located in this field. You can obtain the sector number of the block by using a resident conversion routine, IECOSCR1. The address of this routine is in field CVTOSCR1 of the CVT, and the address of the CVT is in location 16. The routine should be invoked via a BALR 14,15

instruction. If you are passing the track balance to the routine, invoke the routine using a BAL 14,8(15). If you are computing the sector value on modulo devices (3375 and 3380) with variable length records, you must pass the track balance to the sector convert routine.

For RPS devices, the conversion routine does all its work in general registers. You must load registers 0, 2, 14, and 15 with input to the routine. Register usage is as follows:

Register Use

0 For fixed, standard blocks or fixed, unblocked records not in a partitioned data set: Register 0 must be loaded with a 4-byte value in the form XXKR, where XX is a 2-byte field containing the physical block size, K is a 1-byte field containing the key length, and R is a 1-byte field containing the number of the record for which a sector value is desired. The high-order bit of register 0 must be turned off (set to 0) to indicate fixed-length records.

Passing the track balance: Register 0 must be loaded with the 4-byte value of the track balance of the record preceding the required record.

For all other cases: Register 0 must be loaded with a 4-byte value in the form BBIR, where BB is the total number of key and data bytes on the track up to, but not including, the target record; I is a 1-byte key indicator (1 for keyed records, 0 for records without keys); and R is a 1-byte field containing the number of the record for which a sector value is desired. The high-order bit of register 0 must be turned on (set to 1) to indicate variable-length records.

1 Not used by the sector-convert routine.

2 Must be loaded with a 4-byte field where the first byte is the UCB device type code for the device (obtainable from UCB+19), and the remaining three bytes are the address of a 1-byte area that is to receive the sector value.

3-8 Not used.

9-11 Used by the convert routine and are not saved or restored.

12,13 Not used.

14 Must be loaded with the address in which control is to be returned after execution of the sector conversion routine.

15 Used by the conversion routine as a base register and must be loaded with the address of the entry point to the conversion routine.

CHAPTER 4. PASSWORD PROTECTING DATA SETS

The password protection described in this chapter does not apply to VSAM data sets. Information about VSAM data set protection is in VSAM Administration: Macro Instruction Reference and Access Method Services Reference. For information on RACF and its relationship to password protection, see RACF General Information Manual. To use the data set protection feature of the operating system, you must create and maintain a PASSWORD data set consisting of records that associate the names of the protected data sets with the password assigned to each data set. There are four ways to maintain the PASSWORD data set:

- You can write your own routines.
- You can use the PROTECT macro instruction.
- You can use the utility control statements of the IEHPROGM utility program.
- If you have TSO, you can use the TSO PROTECT command.

This chapter discusses only the first two of the four ways: It provides technical detail about the PASSWORD data set that is necessary for writing your own routines, and it describes how to use the PROTECT macro instruction. (The last two of the four ways are discussed in other publications, as indicated in the list of publications below.)

Before using the information in this chapter, you should be familiar with information in several related publications. The following publications are recommended:

- Data Administration Guide contains a general description of the data set protection feature.
- System Messages contains a description of the operator messages and replies associated with the data set protection feature.
- JCL contains a description of the data definition (DD) statement parameter used to indicate that a data set is to be password protected.
- DADSM and CVAE Diagnosis Guide and DADSM Diagnosis Reference contain
- Utilities contains a description of how to maintain the PASSWORD data set using the utility control statements of the IEHPROGM utility program.
- TSO Command Language Reference describes the use of the TSO PROTECT command.

PROVIDING DATA SET SECURITY

In addition to the usual label protection that prevents the opening of a data set without the correct data set name, the operating system provides data set security options that prevent unauthorized access to confidential data. Password protection prevents access to data sets until a correct password is entered by the system operator, or, for TSO, by a remote terminal operator.

The following are the types of access allowed to password-protected data sets:

- PWREAD/PWRITE—A password is required for read or write.
- PWREAD/NOWRITE—A password is required for read. Writing is not allowed.
- NOPWREAD/PWRITE—Reading is allowed without a password. A password is required to write.

To prepare for use of the data set protection feature of the operating system, you place a sequential data set, named `PASSWORD`, on the system residence volume. This data set must contain at least one record for each data set placed under protection. In turn, each record contains a data set name, a password for that data set, a counter field, a protection mode indicator, and a field for recording any information you desire to log. On the system residence volume, these records are formatted as a "key area" (data set name and password) and a "data area" (counter field, protection mode indicator, and logging field). The data set is searched on the "key area."

Note: The area allocated to the data set should not have been previously used for a `PASSWORD` data set, as this may cause unpredictable results when adding records to the data set.

You can write routines to create and maintain the `PASSWORD` data set. If you use the `PROTECT` macro instruction to maintain the `PASSWORD` data set, see "Maintaining the `PASSWORD` Data Set (`PROTECT` Macro)" on page 88. If you use the `IEHPROGM` utility program to maintain the `PASSWORD` data set, see Utilities. These routines may be placed in your own library or in the system's library (`SYS1.LINKLIB`). You may use a data management access method or `EXCP` programming to read from and write to the `PASSWORD` data set.

If a data set is to be placed under protection, it must have a protection indicator set in its label (format-1 DSCB or header 1 tape label). This is done by the operating system when the data set is created, by the `IEHPROGM` utility program, or by the `PROTECT` macro when creating or adding the control password. The protection indicator is set in response to a value in the `LABEL=` operand of the `DD` statement associated with the data set being placed under protection. The publication JCL describes the `LABEL` operand.

Note: Data sets on magnetic tape are protected only when standard labels are used.

Password-protected data sets can only be accessed by programs that can supply the correct password. When the operating system receives a request to open a protected data set, it first checks to see whether the data set has already been opened for this job step. If so, only the access mode will be checked to determine whether it is compatible with the protection mode under which it was previously opened. If the data set has not been previously opened by this job step or if the access mode is not compatible with the protection mode under which it was previously opened, a message is issued that asks for the password; the message goes to the operator console. If the program requesting that the data set be opened is running under TSO in the foreground, the message goes to the TSO terminal operator. If you want the password supplied by another method in your installation, you can modify the `READPSWD` source module or code a new routine to replace `READPSWD` in `SYS1.LPALIB`.

PASSWORD DATA SET CHARACTERISTICS

The PASSWORD data set must reside on the same volume as your operating system. The space you allocate to the PASSWORD data set must be contiguous, that is, its DSCB must indicate only one extent. The amount of space you allocate depends on the number of data sets your installation wants to protect. Each entry in the PASSWORD data set requires 132 bytes of space. The organization of the PASSWORD data set is physical-sequential; the record format is unblocked, fixed-length records (RECFM=F). Each record that forms the data area is 80 bytes long (LRECL=80,BLKSIZE=80) and is preceded by a 52-byte key (KEYLEN=52). The key area contains the fully qualified data set name of as many as 44 bytes and a password of one to eight bytes, left justified with blanks added to fill the areas. The password assigned may be from one to eight alphameric characters in length. DADSM and CVAF Diagnosis Guide and DADSM Diagnosis Reference describe the PASSWORD data set record format.

Note: For data sets on magnetic tape designed according to the specifications of the International Organization for Standardization (ISO) 1001-1979, the equivalent American National Standards Institute (ANSI) X3.27-1978, or the Federal Information Processing Standards (FIPS) 79, do not include generation and version numbers as part of generation data set names. The generation and version numbers are not included as part of the names in the tape labels and are ignored if included in the PASSWORD data set.

You can protect the PASSWORD data set itself by creating a password record for it when your program initially builds the data set. Thereafter, the PASSWORD data set cannot be opened (except by the operating system routines that scan the data set) unless the operator enters the password.

Note: If a problem occurs on a password protected system data set, maintenance personnel must be provided with the password in order to access the data set and resolve the problem.

CREATING PROTECTED DATA SETS

A data definition (DD) statement parameter (LABEL=) may be used to indicate that a data set is to be password protected. For data sets on DASD, an alternative method is to use the PROTECT macro instruction for a previously allocated data set. A data set may be created and the protection indicator set in its label without entering a password record for it in the PASSWORD data set.

Operating procedures at your installation must ensure that password records for all data sets currently password-protected are entered in the PASSWORD data set. Installations where independent computing systems share common DASD resources must ensure that PASSWORD data sets on all systems contain the appropriate password records for any protected data set on shared DASD.

Under certain circumstances, the order in which data sets are allocated and unallocated from multiple systems on shared DASD may result in loss of password protection. For example, if an unprotected data set is allocated and opened by a user on System A and then scratched by a different user on System B, the first user is given a "window" to the unallocated (free) area. If any data set, protected or unprotected, is allocated in that space by a user on either system during the time the "window" is open, the new data set has no protection from the user with the "window."

While the allocation disposition is still NEW, a password-protected data set can be used without supplying a password. However, once the data set is unallocated, any subsequent attempt to open will result in termination of the program unless the password record is available and the correct

password is supplied. Note that, if the protection mode is NOPWREAD and the request is to open the data set for input or read backward, no password will be required.

Tape Volumes Containing More Than One Password-Protected Data Set

To password protect a data set on a tape volume containing other data sets, you must password protect all the data sets on the volume. (Standard labels—SL, SUL, AL, or AUL—are required. For definitions of these label types and the protection-mode indicators that can be used, see Magnetic Tape Labels and File Structure.)

If you issue an OPEN macro instruction to create a data set following an existing, password-protected data set, the password of the existing data set will be verified during open processing for the new data set. The password supplied must be associated with a PWRITE protection-mode indicator.

PROTECTION FEATURE OPERATING CHARACTERISTICS

The topics that follow provide information concerning actions of the protection feature in relation to termination of processing, volume switching, data set concatenation, SCRATCH and RENAME functions, and counter maintenance.

Termination of Processing

Processing is terminated when:

1. The operator cannot supply the correct password for the protected data set being opened after two tries.
2. A password record does not exist in the PASSWORD data set for the protected data set being opened.
3. The protection-mode indicator in the password record and the method of I/O processing specified in the Open routine do not agree, for example, OUTPUT specified against a read-only protection-mode indicator.
4. There is a mismatch in data set names for a data set involved in a volume switching operation. This is discussed in the next paragraph.

Volume Switching

The system ensures a continuation of password protection when volumes of a multivolume data set are switched. It accepts a newly-mounted tape volume to be used for input or a newly-mounted direct access volume, regardless of its use, if these conditions are met:

- The data set name in the password record for the data set is the same as the data set name in the JFCB. (This ensures that the problem program did not change the data set name in the JFCB since the data set was opened.)
- The protection-mode indicator in the password record is compatible with the processing mode, and a valid password has been supplied.

The system accepts a newly-mounted tape volume to be used for output under any of these conditions:

- The security indicator in the HDR1 label indicates password protection; the data set name in the password record is the same as the data set name in the JFCB; and the protection-mode indicator is compatible with the processing

mode. (If the data set name in the JFCB has been changed, a new password is requested from the operator.)

- The security indicator in the HDR1 label does not indicate password protection. (A new label will be written with the security indicator indicating password protection.)
- Only a volume label exists. (A HDR1 label will be written with the security indicator indicating password protection.)

Data Set Concatenation

A password is requested for every protected data set that is involved in a concatenation of data sets, regardless of whether the other data sets involved are protected or not.

SCRATCH and RENAME Functions

To delete or rename a protected data set, it is necessary that the job step making the request be able to supply the password. The system first checks to see if the job step is currently authorized to write to the data set. If not, message IEC301A is issued to request the password. The password provided must be associated with a "WRITE" protection-mode indicator.

Counter Maintenance

The operating system increments the counter in the password record on each usage, but no overflow indication will be given (overflow after 65535 openings). You must provide a counter maintenance routine to check and, if necessary, reset this counter.

MAINTAINING THE PASSWORD DATA SET (PROTECT MACRO)

To use the PROTECT macro instruction, your PASSWORD data set must be on the system residence volume. The PROTECT macro can be used to:

- Add an entry to the PASSWORD data set.
- Replace an entry in the PASSWORD data set.
- Delete an entry from the PASSWORD data set.
- Provide a list of information about an entry in the PASSWORD data set; this list will contain the security counter, access type, and the 77 bytes of security information in the "data area" of the entry.

In addition, the PROTECT macro updates the DSCB of a protected direct access data set to reflect its protection status; this feature eliminates the need for you to use job control language when you protect a data set.

PASSWORD DATA SET CHARACTERISTICS AND RECORD FORMAT (WITH PROTECT MACRO)

When you use the PROTECT macro, the record format and characteristics of the PASSWORD data set are no different from the record format and characteristics that apply when you use your own routines to maintain it.

Number of Records for Each Protected Data Set

When you use the PROTECT macro, the PASSWORD data set must contain at least one record for each protected data set. The password (the last 8 bytes of the "key area") that you assign when you protect the data set for the first time is called the control password. In addition, you may create as many secondary records for the same protected data set as you need. The passwords assigned to these additional records are called secondary passwords. This feature is helpful if you want several users to have access to the same protected data set, but you also want to control the way they can use it. For example: One user could be assigned a password that allowed the data set to be read and written, and another user could be assigned a password that allowed the data set to be read only.

Note: The PROTECT macro will update the protection-mode indicator in the format-1 DSCB in the protected data set only when you issue it for adding, replacing, or deleting a control password.

Protection-Mode Indicator

You can set the protection-mode indicator (third data byte) in the password record to one of four different values:

- X'00' to indicate that the password is a secondary password and the protected data set is to be read only (PWREAD).
- X'80' to indicate that the password is the control password and the protected data set is to be read only (PWREAD).
- X'01' to indicate that the password is a secondary password and the protected data set is to be read and written (PWREAD/PWRITE).
- X'81' to indicate that the password is the control password and the protected data set is to be read and written (PWREAD/PWRITE).

Because of the sequence in which the protection status of a data set is checked, the following defaults will occur:

If control password is: **Secondary password must be:**

- | | |
|---------------------------------------|------------------------------------|
| 1. PWREAD/PWRITE or
PWREAD/NOWRITE | PWREAD/PWRITE or
PWREAD/NOWRITE |
| 2. NOPWREAD/PWRITE | NOPWREAD/PWRITE |

If the control password is set to either of the settings in item 1 above, the secondary password will be set to PWREAD/PWRITE if you try to set it to NOPWREAD/PWRITE.

If the control password is changed from either of the settings in item 1 to the setting in item 2 above, the secondary password will be automatically reset to NOPWREAD/PWRITE.

If the control password is changed from the setting in item 2 to either of the settings in item 1 above, the secondary password is set by the system to PWREAD/PWRITE.

Because the DSCB of the protected data set is updated only when the control password is changed, you may request protection attributes for secondary passwords that conflict with the protection attributes of the control password.

| PROTECT MACRO SYNTAX

The format is:

[symbol]	PROTECT	parameter list address
----------	---------	------------------------

parameter list address—A-type address, (2-12), or (1) indicates the location of the parameter list. The parameter list must be set up before the PROTECT macro is issued. The address of the parameter list may be passed in register 1, in any of the registers 2 through 12, or as an A-type address. The first byte of the parameter list must be used to identify the function (add, replace, delete, or list) you want to perform. See Figure 18 on page 91 through Figure 21 on page 94 for the parameter lists and codes used to identify the functions.

Note: The parameter lists and the areas addressed by the list must reside below 16-megabyte virtual.

PROTECT MACRO PARAMETER LISTS

0	X'01'	13	Control password pointer
1	00 00 00	16	Number of volumes
4	Data set name length	17	Volume list pointer
5	Data set name pointer	20	Protection code
8	00	21	New password pointer
9	00 00 00	24	String length
12	00	25	String pointer

Notes:

- 0 X'01' Entry code indicating ADD function.
- 4 Data set name length.
- 5 Data set name pointer.
- 13 Control password pointer.
The control password is the password assigned when the data set was placed under protection for the first time. The pointer can be 3 bytes of binary zeros if the new password is the control password.
- 16 Number of volumes.
If the data set is not cataloged and you want to have it flagged as protected, you must specify the number of volumes in this field. A zero indicates that the catalog information should be used.
- 17 Volume list pointer.
If the data set is not cataloged and you want to have it flagged as protected, you provide the address of a list of volume serial numbers in this field. Zeros indicate that the catalog information should be used.
- 20 Protection code.
A one-byte number indicating the type of protection: X'00' indicates default protection (for the ADD function; the default protection is the type of protection specified in the control password record of the data set); X'01' indicates that the data set is to be read and written; X'02' indicates that the data set is to be read only; and X'03' indicates that the data set can be read without a password, but a password is needed to write into it. The PROTECT macro will use the protection code value, specified in the parameter list, to set the protection-mode indicator in the password record.
- 21 New password pointer.
If the data set is being placed under protection for the first time, the new password becomes the control password. If you are adding a secondary entry, the new password is different from the control password.
- 24 String length.
The length of the character string (maximum 77 bytes) that you want to place in the optional information field of the password record. If you don't want to add information, set this field to zero.
- 25 String pointer.
The address of the character string that is going to be put in the optional information field. If you don't want to add additional information, set this field to zero.

Figure 18. Parameter List for Add Function

0	X'02'	13	Control password pointer
1	00 00 00	16	Number of volumes
4	Data set name length	17	Volume list pointer
5	Data set name pointer	20	Protection code
8	00	21	New password pointer
9	Current password pointer	24	String length
12	00	25	String pointer

Notes:

0 X'02' Entry code indicating REPLACE function.

4 Data set name length.

5 Data set name pointer.

9 Pointer to current password.
The address of the password that is going to be replaced.

13 Control password pointer.
The address of the password assigned to the data set when it was first placed under protection. The pointer can be set to 3 bytes of binary zeros if the current password is the control password.

16 Number of volumes.
If the data set is not cataloged and you want to have it flagged as protected, you have to specify the number of volumes in this field. A zero indicates that the catalog information should be used.

17 Volume list pointer.
If the data set is not cataloged and you want to have it flagged as protected, you have to provide the address of a list of volume serial numbers in this field. If this field is zero, the catalog information will be used.

20 Protection code.
A one-byte number indicating the type of protection: X'00' indicates that the protection is default protection (for the REPLACE function the default protection is the protection specified in the current password record of the data set); X'01' indicates that the data set is to be read and written; X'02' indicates that the data set is to be read only; and X'03' indicates that the data set can be read without a password, but a password is needed to write into the data set.

21 New password pointer.
The address of the password that you want to replace the current password.

24 String length.
The length of the character string (maximum 77 bytes) that you want to place in the optional information field of the password record. Set this field to zero if you don't want to add additional information.

25 String pointer.
The address of the character string that is going to be put in the optional information field of the password record. Set the address to zero if you don't want to add additional information.

Figure 19. Parameter List for REPLACE Function

0	X'03'	9	Current password pointer
1	00 00 00	12	00
4	Data set name length	13	Control password pointer
5	Data set name pointer	16	Number of volumes
8	00	17	Volume list pointer

Notes:

0 X'03' Entry code indicating DELETE function.

4 Data set name length.

5 Data set name pointer.

9 Current password pointer.

The address of the password that you want to delete. You can delete either a control entry or a secondary entry.

13 Control password pointer.

The address of the password assigned to the data set when it was placed under protection for the first time. The pointer can be 2 bytes of binary zeros if the current password is also the control password.

16 Number of volumes.

If the data set is not cataloged and you want to have it flagged as protected, you must specify the number of volumes in this field. A zero indicates that the catalog information should be used.

17 Volume list pointer.

If the data set is not cataloged and you want to have it flagged as protected, you must provide the address of a list of volume serial numbers in this field. If this field is zero, the catalog information will be used.

Figure 20. Parameter List for DELETE Function

0	X'04'	5	Data set name pointer
1	80-byte buffer pointer	8	00
4	Data set name length	9	Current password pointer

Notes:

0 X'04' Entry code indicating LIST function.

1 80-byte buffer pointer.

The address of a buffer where the list of information can be returned to your program by the macro instruction.

4 Data set name length.

5 Data set name pointer.

9 Current password pointer.

The address of the password of the record that you want listed.

Figure 21. Parameter List for LIST Function

RETURN CODES FROM THE PROTECT MACRO

When the PROTECT macro finishes processing, register 15 contains one of the following return codes:

Code	Meaning
0(X'00')	The PASSWORD data set was successfully updated.
4(X'04')	The PASSWORD of the data set name was already in the password data set.
8(X'08')	The password of the data set name was not in the PASSWORD data set.
12(X'0C')	A control password is required or the one supplied is incorrect.
16(X'10')	The supplied parameter list was incomplete or incorrect.
20(X'14')	There was an I/O error in the PASSWORD data set.
24(X'18') ¹	The PASSWORD data set was full.
28(X'1C')	The validity check of the buffer address failed.
32(X'20') ²	The LOCATE macro failed. LOCATE's return code is in register 1, and the number of indexes searched is in register 0.
36(X'24') ²	The OBTAIN macro failed. OBTAIN's return code is in register 1.
40(X'28') ²	The DSCB could not be updated.
44(X'2C')	The PASSWORD data set does not exist.
48(X'30') ²	Tape data set cannot be protected.
52(X'32') ²	Data set in use.

Notes:

- ¹ For this return code, a message is written to the console indicating that the PASSWORD data set is full.
- ² For this return code, the PASSWORD data set has been updated, but the DSCB has not been flagged to indicate the protected status of the data set.

CHAPTER 5. EXIT ROUTINES

This chapter discusses how installation-written exit modules can:

- Take control before and after direct access device storage management (DADSM) processing
- Take control during Open for a DCB
- Determine whether a missing data set control block (such as for a data set that has been moved to another volume) can be restored to a volume
- Recover from errors that may occur during the opening, closing, or handling of an end-of-volume condition for a data set associated with the user's task

Notes:

1. For information about other available installation-written exits, see Data Administration Guide.
2. For information on IBM-supplied exits for tapes with International Organization for Standardization (ISO), American National Standard Institute (ANSI), or Federal Information Processing Standards (FIPS) labels, see Magnetic Tape Labels and File Structure.

DADSM PREPROCESSING AND POSTPROCESSING EXIT ROUTINES

There are exit routines to enable an installation-written module to take control before and after DADSM processing. An exit parameter list is used to communicate with DADSM. The format of this parameter list is shown in Figure 22 on page 98.

THE EXIT MODULES

All DADSM functions (allocate, extend, scratch, partial release, and rename) have a common preprocessing exit routine and a common postprocessing exit routine that the installation exit routine can replace. These exit routines enable you to gain control before and after DADSM processing. The preprocessing exit routine module is IGGPRE00; the postprocessing exit routine module is IGGPOST0. Each is used by all the DADSM functions listed above. The modules reside in SYS1.LPALIB. You can use System Modification Program (SMP) to replace the IBM-supplied exit routine modules with an installation exit routine you write.

THE EXIT ENVIRONMENT

The exit routines are given control in supervisor state and protect key zero with no locks held. DADSM assumes that the exits are AMODE=24 and RMODE=24. The exit routines must be reentrant. DADSM or the program that invokes DADSM (by issuing enqueue, reserve, and so forth) will have acquired the system resources needed to serialize system functions. These enqueues may prevent other system services from completing successfully. In particular, exit routines must not issue dynamic allocation, OPEN, CLOSE, EOVS, LOCATE, and other DADSM functions because they issue an enqueue on the SYSZTIOT resource. If the exit routines require access to an installation data set, the control blocks required to access that data set (DCB, DEB) should be built during system initialization (IPL/NIP).

The type and number of resources held by DADSM depend upon the DADSM function and the exit taken. For example, on entry to the installation preprocessing exit (IGGPREE0), DADSM holds an enqueue on the VTOC and a reserve on the device for the subject volume of a SCRATCH, RENAME, or partial release function. DADSM releases these resources before the installation postprocessing exit (IGGPOST0) takes control.

You must anticipate system resource contention when services are requested from an exit routine. For example, RACF services issue an enqueue on the RACF data set or a reserve on that data set's volume. This contention can cause system performance problems or an interlock condition.

WHEN IGGPRE00 GETS CONTROL

The preprocessing exit routine, IGGPRE00, is given control before the first VTOC update and after the initial validity check is successful. Input to IGGPRE00 is a parameter list, mapped by macro IECIEXPL, that contains addresses of input data and a function code that identifies the DADSM function. IGGPRE00 is given control once for each volume in the volume list supplied to scratch and rename. A field in the parameter list, IEXRSVWD, may be used to pass data from the preprocessing exit routine to the postprocessing exit routine.

A zero return code from IGGPRE00 indicates the DADSM function may proceed.

REJECTING A DADSM REQUEST

A preprocessing exit routine may reject a DADSM request, in which case an I/O error return code is generated for all functions except allocate and extend. A return code of 4 or 8 from IGGPRE00 to allocate will cause allocate to return X'B4' or X'B0', respectively, to its caller in register 15. Scheduler allocation will treat a X'B4' as a conditional rejection of the allocate request only for the volume being processed. If the allocate request is not for a specific volume, another volume may be chosen and the allocate function retried. Scheduler allocation will treat a X'B0' return code from allocate as an unconditional rejection of the allocate request. If the allocate request is rejected, the preprocessing exit routine can put a reason code in the parameter list field, IEXREASN, and the code will be returned by allocate to its caller, together with the X'B0' or X'B4' return code in register 15. The reason code will appear in the JCL error message if the allocate request is not retried. A nonzero return code from IGGPRE00 to extend will cause extend to return an error return code of X'FFFF FFEC' to its caller. If the caller is end-of-volume, an E37-0C abend will be issued.

DATA THAT DADSM PASSES TO THE EXITS

The format of the parameter list (IEPL) is shown in Figure 22 on page 98.

Name	Offset	Bytes	Description
IEXID	00(X'00')	4	EBCDIC 'IEPL'
IEXLENG	04(X'04')	1	Length of parameter list
IEXFUNC	05(X'05')	1	DADSM function code:
IEXALL			X'01'-Allocate
IEXEXT			X'02'-Extend
IEXSCR			X'03'-Scratch
IEXPR			X'04'-Partial release
IEXREN			X'05'-Rename
IEXEXTCD	06(X'06')	1	Extend code
			X'01' Extend data set on current volume
			X'02' Extend an OS catalog on current volume
			X'04' Extend data set on new volume
			X'81' Extend VSAM data space on current volume
IEXFLAG	07(X'07')	1	Flag byte
IEXENQ		1... ..	VTOC is enqueued upon entry
IEXVIO		.1... ..	VIO data set
IEXMF1		..1.	IEXFMT1 points to DX1FMTID of a partial format-1 DSCB (partial DSCB passed as input to allocate, and not JFCB is not available).
*		...x xxxxx	Reserved
IEXREASN	08(X'08')	2	Installation reject reason code
*		2	Reserved
IEXUCB	12(X'0C')	4	Address of UCB. The UCB address is not available to the pre-exit for VIO allocation.
IEXPTR1	16(X'10')	4	Address of the following:
			- JFCB (allocate, extend, partial release)
			- Scratch/rename input parameter list (in user storage)
IEXPTR2	20(14)	4	Address of the following:
			- DSAB list (ISAM allocate)
			- DEB (Extend on old volume)
			- DCB (partial release)
			- Current volume list entry (scratch/rename)
IEXDSN	24(X'18')	4	Address of the data set name
IEXFMT1	28(X'1C')	4	Address of the 96-byte data portion of format-1 DSCB (pre exit for scratch; pre and post exit for partial release and rename; post exit for allocate). May be supplied by pre exit of allocate, and extend on new volume, to serve as a model if IEXMF1 and IEXVIO are zero.
IEXFMT2	32(X'20')	4	Address of format-2 DSCB. (ISAM allocate post exit.)
IEXRSV00	36(X'24')	4	Reserved

Figure 22 (Part 1 of 2). Format of DADSM Preprocessing and Post-processing Exit Parameter List

Name	Offset	Bytes	Description
IEXEXTBL	40(X'28')	4	Address of DADSM table (pre- and post-exit for scratch and partial release; post-exit for allocate and extend). For VIO allocate post-exit, this is the address of DS1EXT1 in the virtual FMI DSCB.
IEXDCC	44(X'2C')	4	DADSM completion code (post exit)
IEXRSVWD	48(X'30')	4	Reserved word for use by installation exit. Some of the fields passed to the post-exit (such as IEXUCB for VIO allocation) may be successful.

Figure 22 (Part 2 of 2). Format of DADSM Preprocessing and Post-processing Exit Parameter List

PASSING A MODEL FORMAT-1 DSCB

The preprocessing exit for allocate and extend on a new volume may return, in the parameter list field IEXFMT1, the address of the data portion of a model format-1 DSCB, starting with field DS1FMTID. The DSCB will be moved to the allocate or extend work area before building the format-1 DSCB. The only fields that may be nonzero in the area are the DS1REFD (the data-last-referenced field) and fields currently unused. Failure to zero out all fields, except for DS1REFD and all currently unused fields in the model format-1 DSCB, can result in the abnormal termination of the task or lead to unpredictable results. All other fields will be initialized by allocate or extend.

IEXFMT1 may not be supplied by IGGPRE00 for a VIO allocate request (indicated by flag, IEXVIO, set to one), or, if a partial DSCB instead of a JFCB has been supplied to allocate (indicated by flag, IEXMF1, set to one). In the latter case, IEXFMT1 is passed to IGGPRE00 initialized to the address of the DS1FMTID field of the partial format-1 DSCB (supplied to allocate by its caller) in the allocate work area, and DS1REFD may be initialized by IGGPRE00. If extend was successful, IEXFMT1 is zeroed out prior to taking the post-exit, IGGPOST0.

WHEN IGGPOST0 GETS CONTROL

The postprocessing exit module, IGGPOST0, is given control after a DADSM function has been completed or attempted. IGGPOST0 is given control if IGGPRE00 was given control, whether the DADSM function was successful or not. IGGPOST0 is not given control if IGGPRE00 was not given control, or if the DADSM function terminated abnormally. IGGPRE00 may establish a recovery routine, if required, to clean up system resources. The DADSM recovery routine does not give IGGPOST0 control. Input to IGGPOST0 is the same parameter list passed to IGGPRE00. No return codes from IGGPOST0 are defined.

SYSTEM CONTROL BLOCKS

The DADSM installation exit parameter list contains the address of system control blocks. The mapping macros of those control blocks are listed below, together with the name of the system library on which they reside. One of the macros, ICVARXNT, is only supplied with the optional material.

There is no mapping macro for the SCRATCH/RENAME parameter list or the associated volume list.

For extend and partial release, the address of the JFCB passed to the user exit points to a copy of the real JFCB. Updating the copied JFCB will not result in a corresponding change to the real JFCB.

During EXTEND of a VSAM data set, the exit is passed the address of a dummy DEB. This DEB does not contain any EXTENT information.

REGISTERS AT ENTRY TO DADSM EXITS

At entry to your exit routine, register contents are as follows:

Register	Contents
1	Address of the exit parameter list
13	Address of an 18-word save area
14	Return address to DADSM
15	Address of your exit routine

REGISTERS AT RETURN FROM DADSM EXITS

When you return to DADSM, register contents must be as follows:

Register	Contents
0-14	Same as on entry to your exit routine
15	A return code from IGGPRE00

| RETURN CODES FROM DADSM EXITS

No return codes are defined for IGGPOST0. The IGGPRE00 return codes and their meanings are as follows:

Code	Meaning
00(X'00')	Indicates that you want the DADSM request to be processed
04(X'04')	Indicates that no DADSM request for the current volume is to be processed
08(X'08')	Indicates that you do not want the DADSM request to be processed

DCB OPEN INSTALLATION EXIT

There is an exit that enables an installation-written module to take control during Open for a DCB. An exit parameter list is used by open processing to communicate with the exit module. The format of the parameter list is shown in Figure 23 on page 102.

THE EXIT MODULE

OPEN has an exit module that the installation can replace. This module is IFGOEX0B that resides in load module IGC0001I. The load module resides in SYS1.LPALIB. You can use System Modification Program (SMP) to replace the IBM-supplied exit module with an installation exit you write.

THE EXIT ENVIRONMENT

IFGOEX0B is given control in supervisor state and protect key zero with no locks held. System enqueues will have been issued to serialize system functions. These enqueues may prevent other system services from being invoked. In particular, dynamic allocation, OPEN, CLOSE, EOVS, and DADSM functions should not be invoked because of an enqueue on the SYSZTIOT resource. If the exit requires access to an installation data set, the control blocks required to access that data set (DCB, DEB) should be built during system initialization (IPL/NIP). RACF macros may be invoked from the exit.

OPEN PROCESSING BEFORE THE DCB OPEN EXIT GETS CONTROL

The exit module, IFGOEX0B, is given control whenever OPEN processes a DCB. The exit is taken after the following functions have been performed for the DCB.

- DASD data sets
 - Volume mounted
 - Format-1, -2, and -3 DSCBs read
 - Forward merge from format-1 DSCB to JFCB
- Tape data sets
 - Volume mounted
 - Header labels verified
 - Forward merge from header labels to JFCB
- All data sets
 - Forward merge from JFCB to DCB
 - User DCB OPEN installation exit (if any) taken
 - RACF or password verification processing

OPEN PROCESSING AFTER THE DCB OPEN EXIT GETS CONTROL

The following functions have not yet been performed at the time the exit is given control for the DCB.

- Reverse merge from DCB to JFCB (not all fields are merged)
- Reverse merge from JFCB to format-1 DSCB for DASD data sets (not all fields are merged)
- Header labels written (for output tape data set)
- Access-method-dependent processing (obtain buffers, GETMAIN, and build IOBs and DEB)
- Write JFCB
- Write format-1 DSCB

GETTING CONTROL FROM OPEN

The exit is given control for each DCB being opened, even when two or more DCBs are being opened in parallel with one invocation of OPEN.

The exit is given control from OPEN (SVC 19) and OPEN TYPE=J (SVC 22). The exit is given control from end-of-volume (EOV; SVC 55) and from force-end-of-volume (FEOV; SVC 31) when a concatenation of two sequential data sets with unlike attributes is being processed. In this case, EOV gives control to CLOSE, which gives control to OPEN. The exit is not given control from EOV when a concatenation of two sequential data sets with like attributes is being processed. In this case, EOV does not give control to CLOSE and OPEN. A request by the user program for concatenation with unlike attributes is shown in the DCB by flag DCBOFPPC (bit 4; mask X'08') in field DCBOFLGS being set to one.

DATA THAT OPEN PASSES TO THE EXIT

The parameter list mapped by macro IECEOIEXL is supplied to the installation exit. It contains data and the addresses of control blocks that may be of interest to the exit.

The format of the parameter list is shown in Figure 23 on page 102.

Name	Offset	Bytes	Description
OIEXL	00(00)	0	DCB Open installation exit parameter list
OIEXOOPT	00(00)	1	Open option (last 4 bits).
OIEXRSVD		1111	X'F0' first 4 bits reserved.
OIEX00UT	 1111	15 output
OIEX00IN	111	7 outin
OIEXOUPD	1..	4 update
OIEXOINO	11	3 inout
OIEXORDB	1	1 read backward
OIEXOINP		0 input
OIEXUKEY	01(01)	1	User protect key. Key of user DCB.
OIEXLTH	02(02)	2	Length of OIEXL
OIEXUDCB	04(04)	4	Address of user DCB in user protect key (OIEXUKEY)
OIEXPDCB	08(08)	4	Address of protected copy of DCB used by OPEN
OIEXJFCB	12(0C)	4	Address of JFCB
OIEXDSCB	16(10)	4	Address of data portion of format-1 DSCB
OIEXTIOT	20(14)	4	Address of TIOT entry
OIEXUCB	24(18)	4	Address of UCB

Figure 23. Format of DCB OPEN Installation Exit Parameter List (OIEXL)

Note that two DCB addresses are supplied. OPEN maintains a protected copy of your DCB. You can use OPEN's copy of the DCB to test the DCB fields. If you modify your copy of the DCB, OPEN updates its protected copy when it regains control from the exit. The protect key of your DCB is supplied in the exit parameter list. You must use this key to either get information from or modify your DCB.

Be sure you determine the type of DCB and device passed to the exit before testing access-method or device-dependent fields in the DCB. The sample exit shown in Appendix D, "Example of an OPEN Installation Exit Module" on page 225 gives an example of isolating a QSAM DCB being opened to a DASD or tape device.

The JFCB address supplied to the exit points to a copy of the JFCB that is in the OPEN work area. There may be other JFCBs

associated with the OPEN if ISAM or concatenated partitioned data sets are being opened.

In the case of BDAM, ISAM, and concatenated partitioned data sets, the UCB, whose address is supplied to the exit, may not be the only UCB associated with the DCB being opened. The UCB should not be modified.

The TIOT address supplied is of a TIOT entry (TIOENTRY label in the IEFTTIOT1 macro). In the cases of ISAM and concatenated partitioned data sets, other TIOT entries may be associated with the DCB being opened. If concatenation of unlike attributes is being processed, the TIOT entry may have a blank DDNAME field.

The format-1 DSCB passed to the exit is in the OPEN work area. The address is that of the field DS1FMTID. There may be format-2 and -3 DSCBs associated with the format-1 DSCB. There may be other format-1 through -3 DSCBs associated with the DCB being opened in the cases of ISAM, BDAM, and concatenated partitioned data sets. If the OPEN is to the VTOC, a format-4 DSCB address is passed to the exit; this can be determined by testing field DS1FMTID for a value of X'F4', or the data set name in the JFCBDSNM field of 44X'04'.

| DEFAULTING THE DCB BUFFER NUMBER

If a value has not yet been supplied, the exit may be used to supply an installation-determined value for DCBBUFNO (number of buffers) for QSAM DCBs.

A sample exit program that does this is shown in Appendix D, "Example of an OPEN Installation Exit Module" on page 225.

You should not override a nonzero value in DCBBUFNO for QSAM DCBs without knowing what dependency the user program has on that value. When a buffer pool control block address is in the DCB field DCBBUFCA, you cannot override DCBBUFNO; this indicates that buffers have been acquired before OPEN. If no buffer pool control block address exists, DCBBUFCA is set to one (not zero).

You should not override a zero value in DCBBUFNO for BSAM DCBs when DCBBUFCA is set to one without knowing what dependency the user program has on these values. If the user program does not want OPEN to acquire buffer storage space, it indicates this by setting DCBBUFNO to zero and DCBBUFCA to one. If the user program wants OPEN to acquire buffer storage space, it can override DCBBUFNO with a nonzero value. The user program is then responsible for freeing that space after closing the DCB.

MODIFYING THE JFCB

Whenever the JFCB is modified, code 4 should be returned to OPEN. This will cause OPEN to rewrite the JFCB. The JFCB should not be modified if the user program has set JFCNWRIT (bit 4) in byte JFCBTSDM because it indicates the JFCB should not be written.

A sample exit program that modifies the JFCB is shown in Appendix D, "Example of an OPEN Installation Exit Module" on page 225.

Requesting Partial Release

An example of modifying the JFCB in OPEN's work area to request partial release is shown in Appendix D, "Example of an OPEN Installation Exit Module" on page 225. It sets the following bits to 1, indicating a partial release request: JFCRLSE (bits 0 and 1; mask X'CO') in byte JFCBIND1. This should be done only for DASD physical-sequential or partitioned data sets opened for OUTPUT or OUTIN and processed by either (1) EXCP with a 5-word device-dependent section present in the DCB, (2) BSAM, or (3) QSAM.

Care should be taken in modifying the JFCB release bits. For example, a data set that is opened for output many times, writing varying amounts of data each time, may have to extend after each OPEN, resulting in many small extents and, perhaps, reaching the 16-extent limit. This could result in a B37 abend.

Care should also be taken in setting the JFCBSPAC bits to define the space quantity units when the partial release flag, JFCBRLSE, is also set on. A cylinder allocated extent may be released on a track boundary when JFCBSPAC does not indicate cylinder units or average block length units with ROUND specified. This will cause the cylinder boundary extent to become a track boundary extent, thereby losing the performance advantage of cylinder boundary extents. Zeroing the release indicator and increasing secondary allocation quantity (for example, when the data set has extended a large number of times) may prevent such a B37 abend. Setting the release indicator could result in more space being made available to other users sharing the volume.

Updating the Secondary Space Data

The JFCB may also be modified by updating the secondary space data. Byte JFCBCTRI contains the space request type coded in the DD statement or merged from the format-1 DSCB. Field JFCBSQTY contains the amount of secondary space (in either tracks, cylinders, or average block units). Field JFCBPQTY contains the amount of primary space (in either tracks, cylinders, or average block units).

Setting the contiguous bit (JFCONTIG) to zero may prevent an out-of-space abend where there is enough space, but not enough contiguous space, to satisfy a request to extend the data set.

REGISTERS AT ENTRY TO THE DCB OPEN EXIT

At entry to the exit, register contents are as follows:

Register Contents

1	Address of the DCB OPEN installation exit parameter list
13	Address of an 18-word save area
14	Return address to OPEN
15	Address of the entry point to IFGOEX0B

REGISTERS AT RETURN FROM THE DCB OPEN EXIT

When you return to OPEN, register contents must be as follows:

Register	Contents
0-14	Same as on entry to the exit.
15	A return code from IFGOEX0B

RETURN CODES FROM THE DCB OPEN EXIT

The DCB OPEN exit must pass a return code back to OPEN in register 15. The return codes and their meanings are as follows:

Code	Meaning
00(X'00')	Indicates that the JFCB has not been modified
04(X'04')	Indicates that the JFCB has been modified

OPEN/EOV INSTALLATION EXIT FOR FORMAT-1 DSCB NOT FOUND

The function of the format-1 DSCB-not-found installation exit in OPEN and EOV is to determine whether a missing DSCB (such as a data set that has been migrated to another volume) can be restored to the volume. If your exit module restores the DSCB, it indicates this when it returns control to the control program. The exit module, IFGOEX0A, is given control whenever OPEN or EOV fails to find a format-1 DSCB on a volume. There is an IBM-supplied exit module, IFGOEX0A, in SYS1.LPALIB. If you want to use your own exit module, you must replace IFGOEX0A. Your exit module must have an entry point name of IFGOEX0A. If you do not write your own exit module, processing continues normally because the IBM-supplied exit returns a zero return code.

The exit is taken even under conditions under which abnormal termination ordinarily would not occur. Two examples of these conditions follow:

1. When you have specified DISP=MOD and error recovery processing is taking place because the last volume specified in the JFCB does not contain the DSCB, but an earlier volume does. For this case, if your return code from IFGOEX0A is 0 or if your return code is 4 and the DSCB has not been restored, OPEN and EOV search the other volumes for the DSCB after the exit is taken.
2. Another condition occurs during EOV output when space has not yet been allocated on the new volume. Space is allocated after the exit is taken if your return code from IFGOEX0A is 0 or if your return code is 4 and the DSCB has not been restored.

When a DSCB is not found, IFGOEX0A is given control as follows:

- In system protect key 5 (data management key)
- In supervisor state
- The system resource represented by the SYSZTIOT major name is enqueued for shared control. (This ENQ prevents the exit from invoking system functions such as SCRATCH, RENAME, dynamic allocation, or LOCATE.)

DATA THAT OPEN/EOV PASSES TO THE EXIT

The parameter list pointed to by register 1 consists of two fullwords. The first fullword contains the address of the UCB of the volume for which the format-1 DSCB was not found. The second fullword contains the address of the 44-byte data set name, left justified, and padded with blanks. Bit zero of the second fullword is set to one, indicating the last word in the parameter list.

The data set name must not be modified by the exit. The parameter list, save area, and data set name are in protect key 5 virtual storage, which is not fetch protected. IFGOEX0A must be reenterable. All work areas obtained through GETMAIN must be released through FREEMAIN.

REGISTERS AT ENTRY TO THE FORMAT-1 DSCB NOT FOUND EXIT

At entry to your exit routine, register contents are as follows:

Register Contents

0	If X'00', entry was from OPEN (single volume data set). If X'0C', entry was from OPEN (multivolume data set). If X'0F', entry was from EOVS.
1	Address of the parameter list
2-12	Unpredictable
13	Address of an 18-word save area
14	Return address to OPEN/EOVS
15	Address of entry point to IFGOEX0A

REGISTERS AT RETURN FROM THE FORMAT-1 DSCB NOT FOUND EXIT

When you return to OPEN/EOVS, register contents must be as follows:

Register Contents

2-12	Same as on entry to the exit
15	A return code from the exit

RETURN CODES FROM THE FORMAT-1 DSCB NOT FOUND EXIT

The format-1 DSCB not found exit must pass a return code back to OPEN/EOV as follows:

Code	Meaning
00(X'00')	Processing continues normally. This return code is given if the exit does not restore the DSCB. The IBM-supplied exit module always gives return code 0.
04(X'04')	The volume is searched one more time by OPEN or EOV for the DSCB. This return code is given if IFGOEXOA restores the DSCB to the volume. If the DSCB is again not found, IFGOEXOA is not given control and processing continues normally.
08(X'08')	The task is abnormally terminated without attempting to determine if DISP=MOD error recovery or allocation on the new volume should occur. This return code is given if IFGOEXOA encounters an error and you do not want to continue processing.

You should have IFGOEXOA establish its own error recovery environment (for example, through an ESTAE), intercept any indeterminate errors, and return to the control program with return code 8. Problem determination is the responsibility of your exit module. A write-to-programmer (WTO with routing code 11) or a TPUT (if a TSO region) may be used to issue an informative message.

During a parallel OPEN when two or more DCBs are being opened at the same time and two of the DCBs are opening the same data set, the DSCB may be missing. If IFGOEXOA is called for the first of the two DCBs and restores the DSCB, the channel program attempting to read the DSCB for the second DCB may have been executed before the restoration of the DSCB was complete. IFGOEXOA is then called for the second DCB, even though the DSCB has already been restored. Return from IFGOEXOA with a return code 4 is appropriate in this case.

IFGOEXOA is not given control when you are processing a VSAM data set with an ACB; however, it is given control when you are processing a VSAM data space with a DCB. IFGOEXOA is bypassed if the format-4 DSCB is not found on a volume, even if the OPEN is to the VTOC data set name (data set name of 44 bytes of X'04').

| DATA MANAGEMENT ABEND INSTALLATION EXIT

The abend installation exit provides the ability to recover from abnormal conditions that may occur during the opening, closing, or handling of an end-of-volume condition for a non-VSAM data set associated with the user's task.

When an abnormal condition occurs, control passes to the DCB abend user exit routine if one is provided, and processing continues as specified in the DCB abend user exit routine. (The DCB abend user exit routine gives you some options regarding the actions you want the system to take when a condition arises that may result in abnormal termination of your task. For additional information about the DCB abend user exit routine, see Data Administration Guide.) However, if the DCB abend user exit routine is not specified, or if it specifies immediate abnormal termination of the task, the system passes control to the abend installation exit. If a DCB abend user exit routine is not

provided, control immediately passes to the abend installation exit.

IBM supplies an installation exit module, IFG0199I in SYS1.LPALIB, that handles abend situations caused by tape positioning errors. IFG0199I allows you to retry tape positioning when you receive a system completion code 613 with return code 08 or 0C. To perform recovery actions for data management abend situations (other than those caused by tape positioning errors), you can replace installation exit module IFG0199I by modifying the source code supplied in SYS1.SAMPLIB. IFG0199I receives control in protection key zero, supervisor state. IFG0199I checks the system completion code and the return code to determine whether the abend situation is the result of a tape positioning error. If the system completion code is other than 613 with return code 08 or 0C, control returns to the calling module with return code 0, indicating that the abend should continue. Otherwise, IFG0199I checks the counter in the 4-byte work area to determine whether one attempt to reposition the tape has been made. If no attempt to reposition the tape has been made, IFG0199I issues a return code of 4, indicating that positioning should be retried. If one attempt to reposition the tape has been made, IFG0199I issues message IEC613A to the operator to determine whether to attempt repositioning. If the operator specifies that tape positioning is to be attempted again, a return code of 4 is set, indicating that OPEN is to rewind the tape and attempt positioning. If the operator specifies that tape positioning is not to be retried, control is returned to the calling module with a 0 return code.

DATA THAT OPEN/EOV PASSES TO THE EXIT

The format of the parameter list (OAIXL) is shown in Figure 24.

Word Boundary

	User Prot Key	Option Flags	Reserved	Reserved
+0(00)				
+4(04)	Address of the protected copy of the DCB			
+8(08)	Address of the user's DCB related to the abend			
+12(0C)	Address of the UCB related to the abend			
+16(10)	Address of the JFCB related to the abend			
+20(14)	Address of the TIOT related to the abend			
+24(18)	Abend code - Example X'613000C'			
+28(1C)	4-byte installation work area			

1(01) Option flags:

Bits

0 indicates whether the DCB abend user exit was taken

On exit was taken
Off exit was not taken

1 indicates whether to rewind the tape volume

On rewind the tape volume
Off do not rewind the tape volume

REGISTERS AT ENTRY TO THE DATA MANAGEMENT ABEND EXIT

At entry to the exit routine, register contents are as follows:

Register Contents

1	Address of the parameter list (OAIXL)
13	Address of an 18-word save area
14	Return address to OPEN/EOV
15	Address of the entry point to IFG0199I

REGISTERS AT RETURN FROM THE DATA MANAGEMENT ABEND EXIT

When you return to OPEN/EOV, register contents must be as follows:

Register Contents

2-12	Same as on entry to the exit
15	A return code from the exit

RETURN CODES FROM THE DATA MANAGEMENT ABEND EXIT

The data management ABEND exit must pass a return code back to OPEN/EOV as follows:

Code	Meaning
------	---------

00(X'00')	Continue with the abend in process.
-----------	-------------------------------------

04(X'04')	If the bit 1 option flag is on, rewind the tape volume, set the UCBFSCCT and UCBFSEQ fields in the UCB to zero, and retry the abend in process.
-----------	---

If the bit 1 option flag is off, retry the abend in process.

For abend codes that the installation is allowed to retry, see the DCB Abend Exit section in Chapter 7 of the Data Administration Guide.

Modifying the IBM-Supplied Installation Exit Module: Because the IBM-supplied installation exit module handles only a particular abend situation, you may want to modify the source code of that module to perform corrective actions for other abend situations.

You can obtain a copy of the source code from SYS1.SAMPLIB for modification, using the editing function that is available to you. After you have modified the source code, link-edit it into SYS1.LPALIB. The source program is written in Assembler language and uses only macros in SYS1.MACLIB. If you replace the supplied installation module, the exit module that you supply must have the entry point name IFG0199I and it must be reenterable.

CHAPTER 6. SYSTEM MACRO INSTRUCTIONS

This chapter describes miscellaneous macro instructions that allow you to:

- Modify control blocks (RDJFCB macro)
- Obtain information from control blocks and system tables (DEVTYPE macro)
- Perform track capacity calculations (TRKCALC macro)
- Allocate a data set based on a partial DSCB (REALLOC macro)
- Load a message display on an IBM 3480 Magnetic Tape Subsystem (MSGDISP macro)

Before reading this chapter, you should be familiar with the following publications:

- Assembler H Version 2 Application Programming: Guide contains the information necessary to code programs in the assembler language.
- Debugging Handbook contains format and field descriptions of the data areas referred to in this chapter.

INTRODUCTION

The system macro instructions are described in these functional groupings:

- Mapping (IEFUCBOB, IEFJFCBN, and CVT)
- Obtaining device characteristics (DEVTYPE)
- Manipulating the JFCB (RDJFCB)
- Data security (DEBCHK)
- Manipulating queues (PURGE and RESTORE)
- Performing track capacity calculations (TRKCALC)
- Allocating a DASD data set (REALLOC)
- Loading a message display on an IBM 3480 Magnetic Tape Subsystem (MSGDISP)

MAPPING SYSTEM DATA AREAS

The IEFUCBOB, IEFJFCBN, and CVT macro instructions are used as DSECT expansions that define the symbolic names of fields within the unit control block (UCB), job file control block (JFCB), and communication vector table (CVT), respectively.

The CVT, IEFUCBOB, and IEFJFCBN macro definitions are in a distribution library named SYS1.AMODGEN. Before you can issue the macros, you must copy them from SYS1.AMODGEN into SYS1.MACLIB (the IEBCOPY utility can be used to copy the macros), or SYS1.AMODGEN may be concatenated to the macro library before reference is made to SYS1.AMODGEN.

The fields in these blocks are shown and described in Debugging Handbook.

IEFUCBOB—MAPPING THE UCB

This macro instruction defines the symbolic names of the fields in the unit control block (UCB). The macro does not include a DSECT statement. However, if you specify PREFIX=YES, the DSECT statement is provided.

The format is:

<code>[symbol]</code>	<code>IEFUCBOB</code>	<code>[LIST={NO YES}]</code> <code>[,PREFIX={NO YES}]</code>
-----------------------	-----------------------	---

LIST={NO|YES}

NO

specifies that only the UCB prolog is to be printed.

YES

specifies that the UCB prolog and the rest of the UCB are to be printed.

PREFIX={NO|YES}

NO

specifies that no prefix is to be printed.

YES

specifies that the prefix and main body of the UCB are to be printed. A DSECT statement is included if you specify PREFIX=YES.

IEFJFCBN—MAPPING THE JFCB

This macro instruction defines the symbolic names of the fields in the job file control block (JFCB). The macro does not include a DSECT statement. If you require one, code a DSECT statement before the macro statement.

The format is:

<code>[symbol]</code>	<code>IEFJFCBN</code>	<code>[LIST={NO YES}]</code>
-----------------------	-----------------------	------------------------------

LIST={NO|YES}

NO

specifies that only the JFCB prolog is to be printed.

YES

specifies that the JFCB prolog and the rest of the JFCB are to be printed.

CVT—MAPPING THE CVT

This macro instruction defines the symbolic names of all fields in the communication vector table (CVT).

The format is:

<code>[symbol]</code>	<code>CVT</code>	<code>[DSECT={NO YES}]</code> <code>[,LIST={NO YES}]</code>
-----------------------	------------------	--

DSECT={NO|YES}

NO specifies that you do not want a DSECT.

YES specifies that you want a DSECT.

LIST={NO|YES}

NO specifies that only the CVT prolog is to be printed.

YES specifies that the CVT prolog and the rest of the CVT are to be printed.

OBTAINING I/O DEVICE CHARACTERISTICS

Use the DEVTYPE macro instruction to request information relating to the characteristics of an I/O device, and to cause this information to be placed into a specified area. (The results of a DEVTYPE macro instruction executed before a checkpoint is taken should not be considered valid after a checkpoint/restart occurs.) The IHADVA macro maps the data returned by the DEVTYPE macro.

The topics that follow discuss the DEVTYPE macro, device characteristics, and particular output for particular devices.

DEVTYPE MACRO SPECIFICATION

The format is:

[symbol]	DEVTYPE	<u>ddloc-addrx</u> , <u>area-addrx</u> [,DEVTAB] [,RPS]
----------	---------	--

ddloc-addrx

the name of an 8-byte field that contains the symbolic name of the DD statement to which the device is assigned. The name must be left justified in the 8-byte field, and must be followed by blanks if the name is less than eight characters. The doubleword need not be on a doubleword boundary.

area-addrx

the name of an area into which the device information is to be placed. The area can be two, five, or six fullwords, depending on whether or not the DEVTAB and RPS operands are specified. The area must be on a fullword boundary.

DEVTAB

This operand is only required for direct access devices. If DEVTAB is specified, the following number of words of information is placed in your area:

- For direct access devices: 5 words
- For non-direct access devices: 2 words

If you do not code DEVTAB, one word of information is placed in your area if the reference is to a graphics or teleprocessing device; for any other type of device, two words of information are placed in your area.

RPS

If RPS is specified, DEVTAB must also be specified. The RPS parameter causes one additional full word of RPS information to be included with the DEVTAB information.

Note: Any reference for a DUMMY data set in the DEVTYPE macro instruction will cause eight bytes of zeros to be placed in the output area. Any reference to a SYSIN or SYSOUT data set causes X'00000102' to be placed in word 0 and 32760 (X'00007FF8') to be placed in word 1 in the output area. Any reference to a file allocated to a TSO terminal causes X'00000101' to be placed in word 0 and 32760 (X'00007FF8') to be placed in word 1 in the output area.

| DEVICE CHARACTERISTICS INFORMATION

The following information is placed into your area as a result of issuing a DEVTYPE macro:

Word 0

Describes the device as defined in the UCBTYP field of the UCB. For a complete description of this field, refer to Debugging Handbook.

Word 1

Maximum block size. For direct access devices, this value is the smaller of either the maximum size of a nonkeyed block or the maximum block size allowed by the operating system; for magnetic tape devices, this value is the maximum block size allowed by the operating system. For all other devices, this value is the maximum block size accepted by the device.

If DEVTAB is specified, the next three fullwords contain the following information about direct access devices:

Word 2

Bytes 0-1 The number of physical cylinders on the device, including alternates.

Bytes 2-3 The number of tracks per cylinder.

Note: Before you use bytes 2 and 3, read the description of word 4, byte 1, bit 0.

Word 3

Bytes 0-1 Maximum track length. Note that for the IBM 3375 and 3380 direct access devices, this value is not equal to the value in word 1 (maximum block size) as it is for other IBM direct access devices.

Note: Before using bytes 2 and 3, read the description of word 4.

Byte 2 Block overhead, keyed block—the number of bytes required for gaps and check bits for each keyed block other than the last block on a track.

Byte 3 Block overhead—the number of bytes required for gaps and check bits for a keyed block that is the last block on a track.

Bytes 2-3 Block overhead—the number of bytes required for gaps and check bits for any keyed block on a track including the last block. Use of this form is indicated by a 1 in bit 4, byte 1 of word 4.

Basic overhead—the number of bytes required for the count field. Use of this form is indicated by a one in bit 3, byte 1 of word 4.

Word 4

Byte 0 Block overhead, block without key—the number of bytes to be subtracted from word 3, bytes 2 or 3 or bytes 2 and 3, if a block is not keyed.

If bit 3, byte 1 of word 4 is 1, this byte contains the modulo factor for a modulo device.

Byte 1

Bit 0 If on, the number of cylinders, as indicated in word 2, bytes 0 and 1 are invalid.

Bit 1 Reserved.

Bits 2-3 If on, indicates a 3380 is attached to a 3880 Model 13 or 23.

Bit 3 If on, indicates a modulo device (3375, 3380). To calculate the number of data bytes required for a data block for a modulo device, see the device data in Data Administration Guide.

Bit 4 If on, bytes 2 and 3 of word 3 contain a halfword giving the block overhead for any block on a track, including the last block.

Bits 5-6 Reserved.

Bit 7 If on, a tolerance factor must be applied to all blocks except the last block on the track.

Bytes 2-3 Tolerance factor—this factor is used to calculate the effective length of a block. The calculation should be performed as follows:

Step 1 add the block's key length to the block's data length.

Step 2 test bit 7 of byte 1 of word 4. If bit 7 is 0, perform step 3. If bit 7 is 1, multiply the sum computed in step 1 by the tolerance factor. Shift the result of the multiplication 9 bits to the right.

Step 3 add the appropriate block overhead to the value obtained above.

If bit 3, byte 1 of word 4 is 1, bytes 2 and 3 contain the overhead for the data or key field.

If DEVTAB and RPS are specified, the next fullword contains the following information:

Word 5

Bytes 0-1 R0 overhead for sector calculations

Byte 2 Number of sectors for the device

Byte 3 Number of data sectors for the device

Figure 24 shows the output for each device type that results from issuing the DEVTYPE macro.

RETURN CODES FROM THE DEVTYPE MACRO

Control is returned to your program at the next executable instruction following the DEVTYPE macro instruction. Register 15 contains a return code from the DEVTYPE macro. The return codes and their meanings are as follows:

Code Meaning

00(X'00') Indicates that the information concerning the ddname you specified has been successfully moved to your work area.

04(X'04') Indicates that the ddname you specified was not found.

Device ^{1,2}	Maximum Record Size (Word 1, in Decimal)	DEVTAB (Words 2, 3, and 4, in Hexadecimal)	RPS (Word 5, in Hexadecimal)
2540 Reader	80	Not Applicable	Not Applicable
2540 Reader w/CI	80	Not Applicable	Not Applicable
2540 Punch	80	Not Applicable	Not Applicable
2540 Punch w/CI	80	Not Applicable	Not Applicable
2501 Reader	80	Not Applicable	Not Applicable
2501 Reader w/CI	80	Not Applicable	Not Applicable
3890 Document Processor	80	Not Applicable	Not Applicable
3505 Reader	80	Not Applicable	Not Applicable
3505 Reader w/CI	80	Not Applicable	Not Applicable
3525 Punch	80	Not Applicable	Not Applicable
3525 Punch w/CI	80	Not Applicable	Not Applicable
1403 Printer	120 ³	Not Applicable	Not Applicable
1403 w/UCS	120 ³	Not Applicable	Not Applicable
3203 Model 5 Printer	132	Not Applicable	Not Applicable
3211 Printer	132 ³	Not Applicable	Not Applicable
3262 Model 5 Printer	132	Not Applicable	Not Applicable
4245 Printer	132	Not Applicable	Not Applicable
4248 Printer	132 ⁴	Not Applicable	Not Applicable
3800 Printing Subsystem	136 ⁵	Not Applicable	Not Applicable
3400 (9-track, p.e.)	32760	Not Applicable	Not Applicable
3400 (9-track, d.d.)	32760	Not Applicable	Not Applicable

Figure 24 (Part 1 of 2). Output from DEVTYPE Macro

Device ^{1,2}	Maximum Record Size (Word 1, in Decimal)	DEVTAB (Words 2, 3, and 4, in Hexadecimal)	RPS (Word 5, in Hexadecimal)
3400 (7-track)	32760	Not Applicable	Not Applicable
3480 (18-track)	32760	Not Applicable	Not Applicable
2305 Model 2 Fixed-Head Storage	14660	006000083A0A01215B080200	0140B4B1
3330/3333 Disk Storage	13030	019B0013336DBFBF38000200	00ED807C
3330V MSS Virtual Volume	13030	019B0013336DBFBF38000200	00ED807C
3330 Model 11 (or 3333 Model 11) Disk Storage	13030	032F0013336DBFBF38000200	00ED807C
3340 Disk Storage (35 megabytes)	8368	015D000C2157F2F24B000200	0125403D
3340/3344 Disk Storage (70 megabytes)	8368	0230001E4B36010B52080200	0125403D
3350 Disk Storage	19069	0230001E4B36010B52080200	0185807B
3375 Disk Storage	32760	03BF000C8CA000E0201000BF	0340C4BB
3380 Models A04, AA4, and B04 Disk Storage	32760	0376000FBB6001002010010B	04E0DED6
3380 Models A04, AA4, and B04 Disk Storage (attached to a 3880 Model 13 or 23)	32760	0376000FBB6001002030010B	04E0DED6
2250 Model 3 Display Unit		Not Applicable	Not Applicable

Figure 24 (Part 2 of 2). Output from DEVTYPE Macro

Communication Equipment	Record Size
1030,1050,83B3, TWX,2250,S360	Not Applicable
115A,1130	Not Applicable
2780	Not Applicable
2740	Not Applicable

Notes to Figure 24:

¹ CI—card image feature; d.c.—data conversion; d.d.—dual density; p.e.—phase encoding; UCS—universal character set; w/—with.

² Device codes are presented in Debugging Handbook.

- 3 Although certain models can have a larger line size, the minimum line size is assumed.
- 4 The IBM 4248 Printer returns 132 characters even if the 168 Print Position Feature is installed on the device.
- 5 The IBM 3800 Printing Subsystem can print 136 characters per line at 10-pitch, 163 characters per line at 12-pitch, and 204 characters per line at 15-pitch. The machine default is 136 characters per line at 10-pitch.

READING AND MODIFYING A JOB FILE CONTROL BLOCK

To accomplish the functions that are performed as a result of an OPEN macro instruction, the open routine requires access to information that you have supplied in a data definition (DD) statement. This information is stored by the system in a job file control block (JFCB).

In certain applications, you may find it necessary to modify the contents of a JFCB before issuing an OPEN macro instruction. For example, suppose you are adding records to the end of a sequential data set. You might want to add a secondary allocation quantity to allow the existing data set to be extended when the space currently allocated is exhausted. To assist you, the system provides the RDJFCB macro instruction. This macro instruction causes a specified JFCB to be moved from the scheduler work area (SWA), where it is stored, to an area specified in an exit list. (The use of the RDJFCB macro instruction with an exit list is shown under "Example" on page 118. The symbolic names and field descriptions of the JFCB are contained in Debugging Handbook.) When you subsequently issue the OPEN macro instruction, you must indicate, by specifying the TYPE=J operand, that you want to open the data set using the JFCB in the area you specified.

Caution: If you set the bit JFCNWRIT in the field JFCBTSDM to 1 before you issue the OPEN macro instruction, the JFCB is not written back to the SWA at the conclusion of open processing. OPEN TYPE=J normally moves your program's modified copy of the JFCB to the SWA (Scheduler Work Area), replacing the system copy. To ensure that this move is done, your program must set bit zero (0) of the JFCBMASK+4 field to 1. The JFCBMASK format is shown in the Internal Data Areas section of Open/Close/EOV Logic. If the user JFCB, which the system used to open the data set, is not written back to SWA (JFCNWRIT set on), then errors may occur during EOV or close processing.

Some of the modifications that are commonly made to the JFCB include:

- Moving the creation and expiration date fields of the DSCB into the JFCB (see "Using RDJFCB for MSS Virtual Volumes" below).
- Moving the secondary allocation quantity from the DSCB into the JFCB (see "Using RDJFCB for MSS Virtual Volumes" below).
- Moving the DCB fields from the DSCB into the JFCB.
- Adding volume serial numbers to the JFCB (see "Using RDJFCB for MSS Virtual Volumes" and "RDJFCB Security" below).

Volume serial numbers in excess of five are written to the JFCBX (extension) located in the SWA. The JFCBX cannot be modified by user programs.

- Modifying the data set sequence number field in the JFCB.
- Modifying the number-of-volumes field in the JFCB (see "Using RDJFCB for MSS Virtual Volumes" below).

- Setting bit JFCDQDSP in field JFCBFLG3 to invoke the tape volume DEQ at demount facility (see "DEQ at Demount Facility for Tape Volumes," below).
- Modifying the JFCRBIDO field in the JFCB to cause high-speed positioning to a specific data block on a 3480 tape volume.

RDJFCB—READ A JOB FILE CONTROL BLOCK

The RDJFCB macro instruction causes a job file control block (JFCB) to be moved from the SWA (scheduler work area) into an area of your choice as identified via the EXLST parameter of the DCB macro for each data control block specified.

[symbol]	RDJFCB	(<u>dcb-address</u> , [(options)],...)
----------	--------	---

dcb-address,(options)

(same as the dcbaddress, option1, and option2 operands of the OPEN macro instruction, as shown in Data Administration: Macro Instruction Reference).

Although the option operands are not meaningful during the execution of the RDJFCB macro instruction, these operands can appear in the list form of either the RDJFCB or OPEN macro instruction to generate identical parameter lists, which can be referred to with the execute form of either macro instruction.

Example: In Figure 25 on page 119, the macro instruction at EX1 creates a parameter list for two data control blocks: INVEN and MASTER. In creating the list, both data control blocks are assumed to be opened for input; option2 for both blocks is assumed to be DISP. The macro instruction at EX2 moves the system-created JFCBs for INVEN and MASTER from the SWA into the area you specified, thus making the JFCBs available to your problem program for modification. The macro instruction at EX3 modifies the parameter list entry for the data control block named INVEN and indicates, through the TYPE=J operand, that the problem program is supplying the JFCBs for system use.

```

EX1      RDJFCB (INVEN,,MASTER),MF=L
        .
EX2      RDJFCB MF=(E,EX1)
        .
EX3      OPEN (,(RDBACK,LEAVE)),TYPE=J,MF=(E,EX1)
        .
INVEN    DCB      EXLST=LSTA,...
MASTER   DCB      EXLST=LSTB,...
LSTA     DS      0F
        DC      X'07'
        DC      AL3(JFCBAREA)
        .
JFCBAREA DS      0F,176C
        .
LSTB     DS      0F
        .

```

Figure 25. Sample Code Using RDJFCB Macro

Multiple data control block addresses and associated options may be specified in the RDJFCB macro instruction. This facility makes it possible to read several job file control blocks in parallel.

An exit list address must be provided in each data control block specified by an RDJFCB macro instruction. Each exit list must contain an active entry that specifies the virtual storage address of the area into which a JFCB is to be placed. A full discussion of the exit list and its use is contained in Data Administration Guide. The format of the job file control block exit list entry is as follows:

Types of Exit List Entry	Hexadecimal Code (High-Order Byte)	Contents of Exit List Entry (Low-Order Bytes)
Job file control block	07	Address of a 176-byte area to be provided if the RDJFCB or OPEN (TYPE=J) macro instruction is used. This area must begin on a fullword boundary and must be located within the user's region. Also, users running in 31-bit addressing mode must ensure that this area is located below 16-megabyte virtual.

The virtual storage area into which the JFCB is read must be at least 176 bytes long.

The data control block may be open or closed when this macro instruction is executed.

If the JFCB is read successfully for all DCBs in the parameter list, a return code of 0 is placed in register 15. If the JFCB is not read for any of the DCBs because the DDNAME is blank, or a DD statement is not provided, a return code of 4 is placed in register 15.

Warning: The following errors cause the results indicated:

Error	Result
A DD statement has not been provided.	A return code of 4 is placed in register 15.
DDNAME field in DCB is blank.	A write-to-programmer is issued, the request for this DCB is ignored, and a return code of 4 is placed in register 15.
A virtual storage address has not been provided.	Abnormal termination of task.

Note that, if you want to open a VTOC data set to change its contents (that is, open it for OUTPUT, OUTIN, INOUT, UPDAT, OUTINX, or EXTEND), your program must be authorized under the Authorized Program Facility (APF). APF provides security and integrity for your data sets and programs. Details on how you authorize your program are provided in System Programming Library: Supervisor Services and Macro Instructions.

If the RDJFCB routine fails while processing a DCB associated with your RDJFCB request, your task is abnormally terminated. None of the options available through the DCB ABEND exit, as described in Data Administration Guide, is available when a RDJFCB macro instruction is issued.

When using concatenated data sets, the RDJFCB routine modifies only the first JFCB.

USING RDJFCB FOR MSS VIRTUAL VOLUMES: Care must be taken in using RDJFCB if the data set resides on MSS virtual volumes such that:

- The expiration date added does not conflict with other volumes within the specified MSVGP.
- The secondary allocation quantity should be in cylinder increments and be a multiple or submultiple of the primary allocation quantity to avoid fragmentation.
- The number of volumes must not exceed the number available in the specified MSVGP.
- Any volume serial numbers added to the JFCB should exist in the MSVGP.

RDJFCB SECURITY: The volume serial numbers specified in the user-supplied JFCB will be compared with the volume serial numbers in the system JFCB located in the SWA. Each different volume serial number will be enqueued exclusively. The volumes will stay enqueued until the job step terminates, because the close routines will not dequeue the volumes. If the job step already has the volume open, OPEN TYPE=J will continue. If the volume is enqueued by another job step, a 413 abend will occur with a return code of 04.

Some JFCB modifications can compromise the security of existing password-protected data sets. The following modifications are specifically not allowed, unless the program making the modifications is authorized or can supply the password:

- Changing the disposition of a password-protected data set from OLD or MOD to NEW.
- Changing the data set name of one or more of the volume serial numbers when the disposition is NEW.
- Changing the label processing specifications to bypass label processing.

Note: An authorized program is one that is either in supervisor state, executing in one of the system protection keys (Keys 0 through 7), or authorized under the Authorized Program Facility.

RDJFCB USE BY AUTHORIZED PROGRAMS: If you change the data set name in the JFCB, you should do a system enqueue on the major name of "SYSDSN" for the substituted data set name. To use the correct interface with other system functions (for example, partial release), the ENQUEUE macro should include the TCB of the initiator and the length of the data set name (with no trailing blanks). When you complete processing of the data set, you should use the DEQ macro to release the resources.

DEQ AT DEMOUNT FACILITY FOR TAPE VOLUMES

This facility is intended to be used by long-running programs that create an indefinitely long tape data set (such as a log tape). Use of this facility by such a program permits the processed volumes to be allocated to another job for processing (such as data reduction). This processing is otherwise prohibited unless the indefinitely long data set is closed and dynamically unallocated.

You may invoke this facility only through the RDJFCB/OPEN TYPE=J interface by setting bit JFCDQDSP (bit 0) in field JFCBFLG3 (offset 163 or X'A3') to 1. The volume serial of the tape is dequeued when the volume is demounted by OPEN or EOVS with message IEC502E when all the following conditions are present:

- The tape volume is verified for use by OPEN or EOVS.
- JFCDQDSP is set to 1.
- The program is APF authorized (protect key and supervisor/problem state are not relevant).
- The tape volume is to be immediately processed for output. That is, either OPEN verifies the volume and the OPEN option is OUTPUT, OUTIN, or OUTINX; or EOVS verifies the volume and the DCB is opened for OUTPUT, OUTIN, INOUT, or EXTEND, and the last operation against the data set was an output operation (DCBOFLWR is set to 1).

Note that, in order for EOVS to find JFCDQDSP set to 1, the program must not inhibit the rewrite of the JFCB by setting bit 4 of JFCBTSDM to 1.

The tape volume is considered verified after file protect, label type, and density conflicts have been resolved. The volume is dequeued when demounted after this verification, even if further into OPEN or EOVS processing the volume is rejected because of expiration date, security protection, checkpoint data set protection, or an I/O error.

When the volume serial is dequeued, the volume becomes available for allocation to another job. However, because the volume DEQ is performed without unallocating the volume, care must be exercised both by the authorized program and the installation to prevent misuse of the DEQ at demount facility. A discussion of such misuse follows.

1. The authorized program must not close and reopen the data set using the tape volume DEQ at demount facility. If it does, one of the following can occur:

- a. The dequeued volume may be mounted and in use by another job. When the volume is requested for mounting, for the authorized program, the operator is unable to satisfy the mount. Therefore, the operator must either cancel the requesting job, cancel the job using the volume, wait for the requesting job to time out, or wait for the job using the volume to terminate.
- b. The dequeued volume may be allocated to another job but not yet in use. The operator mounts the volume to satisfy the mount request of the authorized job. When the volume is requested for mounting by the other job, the operator is unable to satisfy the mount request, and is faced with the same choices as in a, above.
- c. The dequeued volume may not yet be allocated to another job and the volume is mounted to satisfy the mount request of the authorized job. Another job may allocate the volume and, when the volume is requested for mounting, the situation is the same as in b, above.

It is the responsibility of the installation that permits a program to run with APF authorization to ensure that it does not close and reopen a data set using the DEQ at demount facility.

2. Care should be exercised when an authorized program uses the DEQ at demount facility (data set 1) but processes another tape data set (data set 2). Assume the same volume serial numbers have been coded in the DD statements for data set 1 and data set 2. As the volumes of data set 1 are demounted, they are dequeued even though those volumes may yet be requested for data set 2. All the problems explained in a, b, and c in 1, above, may occur as data set 2 and another job contend for a dequeued volume.

This problem should not occur, given the intended use of the DEQ at demount facility. That is, a long-running application creating an indefinitely long tape data set. This type of application is not normally invoked through batch execution with user-written DD statements.

3. After a volume has been demounted and dequeued because of the DEQ at demount facility, the volume is not automatically rejected by the control program when mounted in response to a specific or nonspecific mount request. Without the use of the facility, the control program can recognize (by the ENQ) that the volume is in use, and reject the volume. Therefore, operations procedures, in effect to prevent incorrect volumes from being mounted, should be reviewed in the light of reduced control program protection from such errors when the DEQ at demount facility is used. Specifically, if a volume is remounted for an authorized program and the volume had been used previously by that authorized program, duplicate volume serial numbers will exist in the JFCB, and the control program will be unable to release the volume during EOV processing.
4. Checkpoint/restart considerations are discussed in Checkpoint/Restart.

OPEN—INITIALIZE DATA CONTROL BLOCK FOR PROCESSING THE JFCB

The OPEN macro instruction initializes one or more data control blocks so that their associated data sets can be processed.

A full explanation of the operands of the OPEN macro instruction, except for the TYPE=J option, is contained in Data Administration: Macro Instruction Reference. The TYPE=J option, because it is used in conjunction with modifying a JFCB, should be used only by the system programmer or only under the system programmer's supervision.

[symbol]	OPEN	(dcb-addr, [(options)], ...) [, TYPE=J]
----------	------	--

TYPE=J

specifies that, for each data control block referred to, you have supplied a job file control block (JFCB) to be used during initialization. A JFCB is an internal representation of information in a DD statement.

During initialization of a data control block, its associated JFCB may be modified with information from the data control block or an existing data set label or with system control information.

The system always creates a job file control block for each DD control statement. The job file control block is placed in the SWA (scheduler work area). Its position, in relation to other JFCBs created for the same job step, is noted in a table in virtual storage.

When the TYPE=J operand is specified, you must also supply a DD statement. However, the amount of information given in the DD statement is at your discretion, because you can modify many fields of the system-created job file control block. If you specify DUMMY on your DD statement, the open routine will ignore the JFCB DSNAME and open the data set as dummy. (See Figure 25 on page 119 for a coding example that modifies a system-created JFCB.)

Note: The DD statement must specify at least:

- Device allocation (refer to JCL for methods of preventing share status)
- A ddname corresponding to the associated data control block DCBDDNAM field

HIGH-SPEED POSITIONING FOR THE IBM 3480 MAGNETIC TAPE SUBSYSTEM

Fast positioning for 3480 tape drives is available when opening a tape data set on an IBM standard-labeled tape for EXTEND (OUTINX, EXTEND, or DISP=MOD). To invoke high-speed positioning your program must modify certain fields in the JFCB and use OPEN TYPE=J to open the data set.

The following procedure should be used to modify the JFCB:

1. Issue RDJFCB to have the system move the JFCB into your work area.
2. Set the JFCPOSID flag in the JFCBFLG3 flag byte to indicate that you are providing a Block ID for a high-speed search.
3. Move the Block ID into the JFCRBIDO field of the JFCB. (Specify the Block ID of the tape mark immediately following the last block of user data in the data set being opened.)
4. Issue OPEN TYPE=J with the modified JFCB.

Once the tape is positioned, OPEN processes the trailer labels for the data set being extended.

If you set the JFCPOSID flag off, OPEN positions the volume normally, as though the high-speed positioning feature were not active.

If you set the JFCPOSID flag on, but do not provide a block ID in the JFCRBIDO field, OPEN positions the volume normally. OPEN then moves the block ID of the first header label record into the JFCRBIDO field.

If you set the JFCPOSID flag on, but the block ID you provide in the JFCRBIDO field does not exist on the tape, OPEN processing fails. OPEN issues error message IECL47I to indicate this condition.

If the JFCPOSID flag is on during CLOSE processing, CLOSE inserts the block ID for the first header label record of the next data set (which may not exist) into the JFCRBIDO field. Therefore, if you deallocate the 3480 device and want to use the current block ID for subsequent processing, you must save the block ID before you CLOSE the data set.

Notes:

1. If you specify dynamic deallocation (with SVC 99, FREE=CLOSE on the DD statement, or the FREE option on the CLOSE macro), then the block ID for the next data set will not be available to your program. This is because dynamic deallocation frees the JFCB.
2. When using high-speed positioning, specify the data set sequence number normally, either explicitly by LABEL=(seqno,SL) on the DD statement, or by default.

ENSURING DATA SECURITY BY VALIDATING THE DATA EXTENT BLOCK

Protecting one user's data from inadvertent or malicious access by an unauthorized user depends on protection of the data extent block (DEB). The DEB is a critical control block because it contains information about the device a data set is mounted on, and describes the location of data sets on direct access device storage volumes. The DEB also contains the address of the appendage vector table (AVT). Using the AVT, an unauthorized user can modify the AVT to give control to a routine in supervisor state to read from and write to data sets to which access would otherwise be denied.

To guarantee protection of the DEB, the DEBCHK macro instruction is provided. The DEBCHK macro instruction can be found in SYS1.MACLIB. The DEBCHK macro is issued by several components of the system control program. For example:

- The open access method executors issue the macro to add the address of a DEB they have built to a list of valid addresses called the DEB table. The DEB validity-checking routine builds and maintains a DEB table for each job step.
- The EXCP processor uses the macro to verify that the DEB passed with each EXCP request is in the DEB table.
- The close component issues the macro to remove a DEB from the DEB table.

If you code a routine that builds a DEB, you must add the address of the DEB you built to the DEB table. If you code a routine that depends on the validity of a DEB that is passed to your routine, you should verify that the DEB passed to your routine has a valid entry in the DEB table and points to your DCB or access method control block (ACB). Use the TYPE=ADD and the TYPE=VERIFY operands of the macro, respectively.

To prevent an asynchronous routine from changing or deleting, or assigning a new DEB to a DCB, you must hold the local lock. In this case, you must use the branch entry to the DEBCHK verify routine.

Additional details about the functions provided by the DEB validity-checking routine and about the contents of the DEB table are available in Open/Close/EOV Logic.

The DEBCHK macro instruction provides four functions:

- Adds the address of a DEB to the DEB table, which is located in protected storage. The DEB table contains the address of every user DEB associated with a given job step. Every system control program component that builds a user DEB must add the address of that DEB to a DEB table.
- Verifies that the DEB table associated with a given job step contains the address of a valid DEB and that the DEB points to the DCB (or ACB). Any system control program component or problem program can use this function to verify that a DEB is valid.
- Deletes the address of a DEB from the DEB table. Any program that deletes a user DEB must, before it deletes the DEB, issue a DEBCHK macro with a TYPE=DELETE operand to delete the address of the DEB from the DEB table. If the DEB validity-checking routine encounters an error while deleting the address from the DEB table, the job step is abnormally terminated.
- Deletes the address of a DEB from the DEB table in the same way as the preceding function, except that, instead of terminating the job step, this function merely returns an error code in register 15. This function is provided to prevent recurring abnormal termination. The format of the DEBCHK and a description of the operands follow:

DEBCHK—MACRO SPECIFICATION

[symbol]	DEBCHK	<u>cbaddr</u> [,TYPE={VERIFY ADD DELETE PURGE}] [,AM={amtype (amaddr) ((amreg))}] [,BRANCH={NO YES}] [,TCBADDR=address] [,KEYADDR=address] [,SAVREG=reg] [,MF=L]
----------	--------	---

cbaddr

for **BRANCH=NO**
 RX-type address, (2-12), or (1)

A control block address passed to the DEBCHK routine. This operand is ignored if MF=L is coded. For verify, add, and delete requests, cbaddr is the address of a DCB or ACB that points to the DEB whose address is either verified to be in the DEB table, added to the DEB table, or deleted from the DEB table. For the purge function, cbaddr is the address of the DEB whose pointer is to be purged from the table: No reference is made to the DCB or ACB.

Note: A spooled DCB's DEB does not point back to the DCB, but to the spooled ACB; in this case, the DEBCHK should be issued against the ACB.

for **BRANCH=YES**
 The A-type address of a 4-byte field, or a register (1) or (3-12), that points to the DCB or ACB containing the DEB to be verified.

TYPE={VERIFY|ADD|DELETE|PURGE}
 indicates the function to be performed. If MF=L is coded, TYPE is ignored. The functions are:

VERIFY

This function is assumed if the TYPE operand is not coded. The control program checks the DEB table to determine whether the DEB pointer is in the table at the location indicated by the DEBTBLOF field of the DEB. The DEB is also checked to verify that DEBDCBAD

points to the DCB (or ACB) passed to DEBCHK. The DEBAMTYP field in the DEB is compared to the AM operand value, if given. The two must be equal. TYPE=VERIFY may be issued in either supervisor or problem state.

ADD

The DEB and the DCB (or ACB) must point to each other before the DEB address can be added to the DEB table. Before the DEB pointer can be added to the table, the DEB itself must be queued on the current TCB DEB chain (the TCBDEB field contains the address of the first DEB in the chain). The DEB address is added to the DEB table at some offset into the table. That offset value is placed in the DEBTBLOF field of the DEB, and the access method type is inserted into the DEBAMTYP field of the DEB. A zero is placed in the DEBAMTYP field if the AM operand is not coded. TYPE=ADD can be issued only in supervisor state.

DELETE

The DEB and the DCB (or ACB) must point to each other before the DEB address can be deleted from the DEB table. TYPE=DELETE can be issued only in supervisor state.

PURGE

The DEB pointer is removed from the DEB table without checking the DCB (or ACB). TYPE=PURGE can be issued only in supervisor state.

AM

specifies an access method value. Each value corresponds to a particular access method type (note that BPAM and SAM have the same values):

Type	Value
TCAMAP	(X'84')
SUBSYS	(X'81')
ISAM	(X'80')
BDAM	(X'40')
SAM	(X'20')
BPAM	(X'20')
TAM	(X'10')
GAM	(X'08')
TCAM	(X'04')
EXCP	(X'02')
VSAM	(X'01')
NONE	(X'00')

The operand can be coded in one of the following three ways, only the first of which is valid for the list form (MF=L) of the instruction.

amtype

refers to the access method: ISAM, BDAM, SAM, BPAM, TAM (which refers to BTAM only), GAM, TCAM, EXCP, or VSAM. TCAMAP identifies a TCAM application-program DEB. SUBSYS identifies a subsystem of the operating system, such as a job entry subsystem. NONE indicates that no access method or subsystem is specified.

(amaddr)

is the RS-type address of the access method value. This format may not be coded when MF=L is used.

((amreg))

is one of the general registers 1 through 14 that contains the access method value in its low-order byte (bit positions 24 through 31). The high-order bytes are not inspected. This form may not be used when MF=L is coded.

The use of `amaddr` and `amreg` should be restricted to those cases where the access method value has been generated previously by the `MF=L` form of `DEBCHK`. If `MF=L` is not coded, the significance of the `AM` operand depends upon the `TYPE`.

If `TYPE` is `ADD` and `AM` is specified, the access method value is inserted in the `DEBAMTYP` field of the `DEB`, and all subsequent `DEBCHK` macros referring to this `DEB` must either specify the same `AM` or omit the operand. When the `AM` operand is omitted for `TYPE=ADD`, a null value (0) is placed in the `DEB` and all subsequent `DEBCHK` macros must omit the `AM` operand.

If `AM` is specified when the `TYPE` is `PURGE`, `DELETE`, or `VERIFY`, the access method value is compared to the value in the `DEBAMTYP` field of the `DEB`. If `AM` is omitted, no comparison is made.

BRANCH={NO|YES}

specifies whether you want to use the branch entry to the `DEBCHK` verify routines.

NO

specifies branch entry is not to be used. The operands `SAVREG`, `TCBADDR`, and `KEYADDR` are ignored.

YES

specifies the branch entry is to be used. `TYPE=VERIFY` must be implicitly or explicitly specified. The operands `TCBADDR` and `KEYADDR` are required. `AM` and `MF` are ignored. Notes for `BRANCH=YES`:

- Registers 1, 2, 10, 11, 14, and 15 must not be used for `SAVREG=`.
- Registers 1, 2, 10, 11, 14, 15, and the register specified for `SAVREG=` must not be used for `cbaddr`, `TCBADDR=`, or `KEYADDR=`.
- The contents of registers 10, 11, and 14 are unpredictable on completion. Also, if you do not specify `SAVREG=`, the contents of register 2 are unpredictable.
- At completion time, register 1 contains the address of the `DEB`, and register 15 contains either 0, 4, or 16 (see below for codes and their meanings).

TCBADDR=address—A-type address or (3-12)

specifies the location or register containing the address of the `TCB` to be used by the `DEBCHK` verify routine. Use this operand only when `BRANCH=YES`.

KEYADDR=address—A-type address or (3-12)

specifies the location, or a register pointing to the location of a field containing the key to be used when accessing the `DCB` (or `ACB`). Use this operand only when `BRANCH=YES`.

SAVREG=reg

specifies the register in which register 2 is to be saved. Use this operand only when `BRANCH=YES`.

MF=L

indicates the list form of the `DEBCHK` macro instruction. When `MF=L` is coded, a parameter list is built consisting of the access method value that corresponds to the `AM` keyword. This value may be referenced by name in another `DEBCHK` macro by coding `AM=(amaddr)`, or it may be inserted into the low-order byte of a register before issuing another `DEBCHK` macro by coding `AM=((amreg))`.

Return Codes from the DEBCHK Macro

If the DEBCHK routine completes successfully, register 15 will be set to 0 and register 1 will contain the address of the DEB when control is returned to your program. Otherwise, register 15 will contain one of the following decimal codes:

Code	Meaning
04(X'04')	Either (a) the DEB table associated with the job step does not exist; or (b) the DEBTBLOF field of the DEB was set to zero or a negative number, or was larger than the DEB table; or (c) register 1 did not contain the same address as the DEB table entry.
08(X'08')	An invalid TYPE was specified. (The DEBCHK routine was entered by a branch, not by the macro.)
12(X'0C')	Your program was not authorized and TYPE was not VERIFY.
16(X'10')	DEBDCBAD did not contain the address of the DCB (or ACB) that was passed to the DEBCHK routine.
20(X'14')	The AM value does not equal the value in the DEBAMTYP field.
24(X'18')	The DEB is not on the DEB chain and TYPE=ADD was specified.
28(X'1C')	TYPE=ADD was specified for a DEB that was already entered in the DEB table.
32(X'20')	The DEB table exceeded the maximum size (32760 bytes) and TYPE=ADD.

PURGING AND RESTORING I/O REQUESTS

The system's purge routines, guided by a parameter list you pass them, perform either a halt or a quiesce operation. In a halt operation, the purge routines stop the processing of specified I/O requests that were initiated with an EXCP macro instruction. In a quiesce operation, the purge routines:

- Allow the completion of I/O requests that were initiated with an EXCP macro instruction and have been passed to the I/O supervisor for execution
- Stop the processing of those requests that have not as yet been passed to the I/O supervisor, but save the IOBs of the requests so that they can be reprocessed (restored) later.

The system's restore routines make it possible to reprocess I/O requests that are quiesced. (Note: Not covered here is the purge and restore processing that takes in I/O requests not initiated by an EXCP macro instruction. If you want to learn the full scope of purge and restore processing, see the I/O supervisor logic section of System Logic Library, Volume 8.)

You can give control to the purge and restore routines in one of two ways: (1) by loading register 1 with the address of the parameter list and issuing specific SVC instructions or (2) by issuing the PURGE and RESTORE macro instructions. If your installation requires the use of macro instructions, you must add the macro definitions to the macro library (SYS1.MACLIB) or place them in a partitioned data set and concatenate this data set to the macro library. The macro definitions, JCL, and utility statements needed to add the macros to your macro library are presented in Figure 26 on page 129, and Figure 27 on page 129. Whether you issue the macro instructions or the SVC instructions, you must first build a parameter list. The SVC instructions are SVC 16 for PURGE and SVC 17 for RESTORE.

PURGE Macro Definition

```
MACRO
&NAME PURGE      &LIST
AIF      ('&LIST' EQ '').E1
&NAME IHBINNRA  &LIST      LOAD REG 1
SVC      16
MEXIT
.E1 IHBERMAC    01,147      LIST ADDR MISSING
MEND
```

Control Statements Required

```
//jobname JOB      {parameter}
//stepname EXEC     PGM=IEBUPDTE,PARM=NEW
//SYSPRINT DD      SYSOUT=A
//SYSUT2   DD      DSNAME=SYS1.MACLIB,DISP=OLD
//SYSIN    DD      *
./ ADD     NAME=PURGE,LIST=ALL
.
.
PURGE macro definition
.
.
./ ENDUP
/*
```

Figure 26. Macro Definition, JCL, and Utility Statements for Adding PURGE Macro to the System Macro Library

RESTORE Macro Definition

```
MACRO
&NAME RESTORE    &LIST
AIF      ('&LIST' EQ '').E1
&NAME IHBINNRA  &LIST      LOAD REG 1
SVC      17          ISSUE SVC FOR RESTORE
MEXIT
.E1 IHBERMAC    01,150      LIST ADDR MISSING
MEND
```

Control Statements Required

```
//jobname JOB      {parameters}
//stepname EXEC     PGM=IEBUPDTE,PARM=NEW
//SYSPRINT DD      SYSOUT=A
//SYSUT2   DD      DSNAME=SYS1.MACLIB,DISP=OLD
//SYSIN    DD      DATA
./ ADD     NAME=RESTORE,LIST=ALL
.
.
RESTORE macro definition
.
.
./ ENDUP
/*
```

Figure 27. Macro Definition, JCL, and Utility Statements for Adding RESTORE Macro to the System Macro Library

PURGE—HALT OR FINISH I/O-REQUEST PROCESSING

The macro instruction used to call the purge routines is coded as follows:

<u>[symbol]</u> <u>address</u>	PURGE	<u>parameter-list</u>
-----------------------------------	-------	-----------------------

parameter list address—RX-type address, (2-12) or (1) address of a parameter list, 12 or 16 bytes long, that you have built on a fullword boundary in your storage. The parameter list address can be specified as an RX-type constant or in registers 2 through 12, or 1.

The format and contents of the parameter list are as follows:

Byte	Contents
0	<p>A byte in which you specify what the purge routines will do. These are the bit settings and their meanings:</p> <p>1... Purge I/O requests to a single data set.</p> <p>0... Either purge I/O requests associated with a TCB or address space, or purge I/O requests to more than one data set.</p> <p>.1... Post ECBs associated with purged I/O requests.</p> <p>..1. Halt I/O-request processing. (Quiesce I/O-request processing, if 0.)</p> <p>...1 Purge related requests only. (Valid only if a data-set purge is requested.)</p> <p>.... 0... Reserved—must be zero.</p> <p>.... .1.. Do not purge the TCB request-block chain of asynchronously scheduled processing.</p> <p>.... ..1. Purge I/O requests associated with a TCB.</p> <p>.... ...1 This is a 16-byte parameter list. Additional purge options are specified in bytes 12 to 15. (If this bit is off, the list is 12 bytes long, and the purge routines do not put a return code in byte 4 of this list or in register 15.)</p>
1,2,3	<p>The address of a DEB if you're purging I/O requests to a single data set. The address of the first DEB in a chain of DEBs if you're purging I/O requests to more than one data set. (The next-to-the-last word of each DEB must point to the next DEB in the chain; the second word of the last DEB must contain zeros.)</p>
4	<p>A byte of zeros. (If bit 7 of byte 0 is on, the purge routines will put a code in this byte: X'7F' if the purge operation is successful; X'40' if it is not successful. If bit 7 of byte 0 is off, then X'7F' appears in this byte.)</p>
5,6,7	<p>The address of the TCB associated with the I/O requests you want purged (but only if you turned on bit 6 of byte 0). May be zeros if the TCB is the one you're running under.</p>
8	<p>Driver 10. (Default value of X'00' implies that EXCP is the owner.)</p>

9,10,11 The address of a word in your storage or the address of the DEBUSPRG field (which is X'11' bytes more than the DEB address in this parameter list). At whichever address you specify, the purge routines store a pointer to the purged I/O restore list, PIRL. In the PIRL is a pointer to the first IOB in the chain of IOBs. The location of the pointer and format of the chain are shown in Figure 28 on page 133.

Note: This field is relevant for quiesce options only.

12 A byte in which you can specify additional purge options. These are the bit settings and their meanings:

Note: The following applies only if bit 7 of byte 0 is set to one.

...1. Purge I/O requests associated with an address space. (You must be in supervisor state.)

...1 Check the validity of all the DEBs associated with the purge operation if this is a data-set purge. Validate this parameter list, whatever the type of purge operation, by ensuring that there are no inconsistencies in the selection of purge options. (If you are in problem state, these actions are taken regardless of the bit setting.)

.... 1... Ensure that I/O requests will be reprocessed (restored) under their original TCB. (If zero, and this byte is meaningful (bit 7 of byte 0 is on), the I/O requests will be reprocessed under the TCB of the program making the restore request.)

.... .0.. Must be zero.

13 A byte of zeros.

14,15 The 2-byte ID of the address space associated with the I/O requests you want purged. (Only meaningful if bit 2 of byte 12 is on.)

Control is returned to your program at the instruction following the PURGE macro instruction.

Return Codes from PURGE Macro

If the purge operation was successful, register 15 will contain zeros. Otherwise, register 15 will contain one of the following return codes:

Code	Meaning
04(X'04')	Your request to purge I/O requests associated with a given TCB was not honored because that TCB did not point to the job step TCB, as it must when the requestor is in problem state.
08(X'08')	Either you requested an address-space purge operation, but were not in supervisor state, or you requested a data-set purge operation, but supplied no data-area address in bytes 1, 2, and 3 of the purge parameter list.
20(X'14')	Another purge request has preempted your request. You may want to reissue your purge request in a time-controlled loop.

Note: Register 15 will contain zeros, regardless of the outcome of the purge operation, if you set bit 7 in byte 0 of the parameter list to zero.

MODIFYING THE IOB CHAIN

Note that, although it is not a recommended procedure, if you want to change the order in which purged I/O requests are restored or prevent a purged request from being restored, you may change the sequence of IOBs in the IOB chain or remove an IOB from the chain. The address of the IOB chain can be obtained from the PIRL (see Figure 28). (The address of the PIRL is at the location pointed to by bytes 9 through 11 of the purge parameter list.)

PIRL

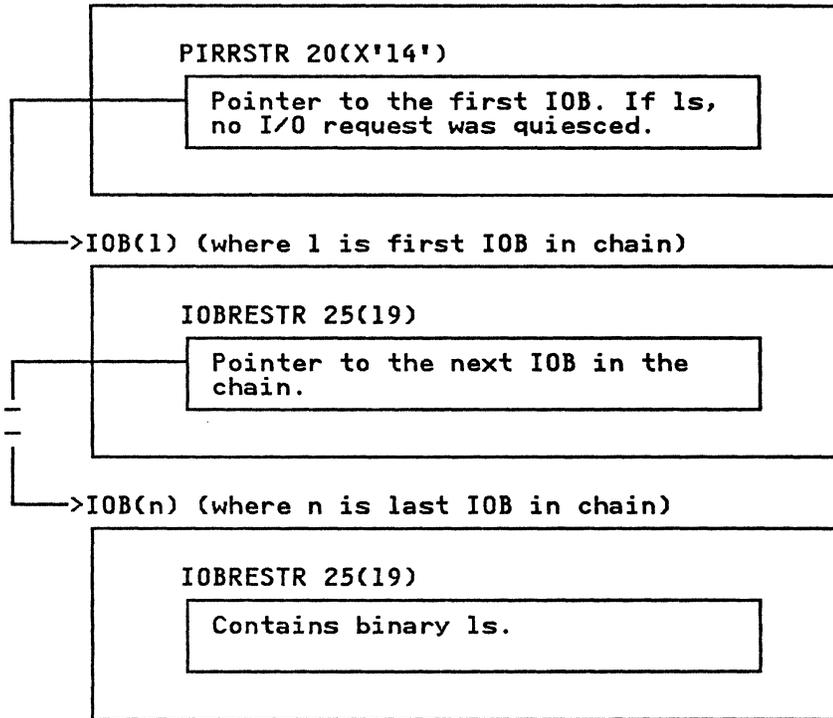


Figure 28. The PIRL and IOB Chain

RESTORE—REPROCESS I/O REQUESTS

The RESTORE macro is coded as follows:

[symbol]	RESTORE	<u>restore address</u>
----------	---------	------------------------

restore address—RX-type address, (2-12) or (1)
address you specified at byte 9 of the purge parameter list.

PERFORMING TRACK CALCULATIONS

The TRKCALC macro performs track capacity calculations. The standard, list, execute, and DSECT forms of the macro are described. Examples of the TRKCALC macro follow the macro descriptions. Using TRKCALC, you may do the following:

- Perform track capacity calculations
- Determine the number of records of a given size that can be written on a full track or on the remainder of a track
- Perform track balance calculations as follows:
 - Determine whether a given record size can be written in the space remaining on the track and return the new track balance.

- Determine the maximum size record that can be written on the track if the given record does not fit.
- Determine the track balance if the last physical record is removed from the track.

TRKCALC—STANDARD FORM

The format of the TRKCALC macro is:

[symbol]	TRKCALC	FUNCTN={TRKBAL TRKCAP} {,DEVTAB=addr ,UCB=addr ,TYPE=addr} [,BALANCE=addr] [,REMOVE={YES NO}] [,MAXSIZE={YES NO}] {,RKDD=addr ,R=addr,K=addr,DD=addr} [,REGSAVE={YES NO}] [,MF=I]
----------	---------	--

FUNCTN={TRKBAL|TRKCAP}

specifies the function to be performed.

Note: You must specify one of the three keywords, DEVTAB, UCB, or TYPE, to provide the macro a source for information.

TRKBAL

if REMOVE=NO is specified, TRKBAL calculates whether an additional record fits on the track and what the new track balance would be if the record were added. If REMOVE=YES is specified, TRKBAL calculates what the track balance would be if a record were removed from the track. The record to be added or removed from the track is defined by the RKDD parameter, or by the R, K, and DD parameters.

If R=1 (or the R value in the RKDD parameter is 1) and REMOVE=NO is specified, record 1 is added to an empty track; if R=1 and REMOVE=YES is specified, record 1 is deleted from the track, leaving an empty track.

If R≠1, the specified record is added to or removed from the track. The input track balance may be supplied through the BALANCE parameter; if it is not supplied, it is assumed that the track contains equal-sized records as specified in the RKDD parameter (or in the R, K, and DD parameters).

When REMOVE=NO is specified, one of the following occurs:

- If the record fits on the track, register 0 contains the new track balance.
- If the record does not fit on the track and MAXSIZE=NO is specified, a "record does not fit" return code is given in register 15.
- If the record does not fit and MAXSIZE=YES is specified, one of the following happens:
 - The data length of the largest record that fits in the remaining space is returned in register 0.
 - A code is returned that indicates no record fits in the remaining space.

When REMOVE=YES is specified, one of the following occurs:

- If R=1, register 0 contains the track capacity.
- If R≠1, registers 0 contains the input track balance (supplied through the BALANCE parameter) incremented by the track balance used by the input record. If the input balance is not supplied, register 0 contains the track capacity left after R-1 records are written on the track.

TRKCAP

calculates, and returns in register 0, the number of fixed-length records that may be written on a whole track (R=1) or on a partially filled track (R≠1). The records are defined by the K and DD values of the RKDD parameter, or by the K and DD parameters.

One of the following occurs:

- If R=1, the BALANCE parameter is ignored and the calculation is made on an empty track.
- If R≠1 and the BALANCE parameter is omitted, the calculation is made for a track that already contains R-1 records of the length defined by the K and DD values.
- If R≠1 and the BALANCE parameter is supplied, the calculation is made for a track whose remaining track balance is the value of the BALANCE parameter.

DEVTAB=addr—RX-type address, (2-12), (0), (14)

addr specifies a word that contains the address of the Device Characteristics Table Entry (DCTE). If you specify a register, it contains the address of the DCTE, not the address of a word containing the address of the DCTE. The address of the DCTE can be found in the DCBDVTBA field of an opened DCB.

UCB=addr—RX-type address, (2-12), (0), (14)

addr specifies the address of a word that contains the address of the UCB. If you specify a register, it contains the address of the UCB, not the address of a word containing the address of the UCB. You must ensure that the UCB address is valid by verifying that byte 3 (UCB+2) in the UCB contains X'FF'.

TYPE=addr—RX-type address, (2-12), (0), (14)

you may specify the address of the UCB device type (UCBTBYT4), or you may specify the 1-byte UCB device type in the low-order byte of a register.

BALANCE=addr—RX-type address, (2-12), (0), (14)

you may specify either the address of a halfword containing the current track balance, or you may specify the balance in the low-order two bytes of a register. The value supplied may be the value returned when you last issued TRKCALC. If R=1, the balance is reset to track capacity by TRKCALC, and your supplied value is ignored. This is an input value and is not modified by the TRKCALC macro. The resulting track balance is returned in register 0 and in the TRKCALC parameter list field STARBAL.

REMOVE={YES|NO}

indicates if a record is to be deleted from the track.

YES

specifies that the record identified by the record number (specified in the R keyword) is being deleted from the track. The track balance is incremented instead of decremented.

Note: YES is valid only on a FUNCTN=TRKBAL call.

NO

specifies that a record is not to be deleted from the track. NO is the default.

MAXSIZE={YES|NO}

YES

If the specified record does not fit, the largest length of a record with the specified key length that fits is returned (register 0).

Note: YES is valid only on a FUNCTN=TRKBAL call.

NO

Maximum size is not returned. NO is the default.

RKDD=addr—RX-type address, (2-12), (0), (14)
addr specifies a word containing a record number (1 byte), keylength (1 byte), and data length (2 bytes) (bytes 0, 1, and 2 and 3, respectively) or a register containing the record number, key length, and data length. R, K, and DD may be specified by this keyword, or you may use the following three keywords instead.

R=addr—RX-type address, (2-12), (0), (14), or n
you may specify either the address of the record number, or you may specify the record number using the low-order byte of a register or immediate data (n). Specify a decimal digit for n (immediate data).

K=addr—RX-type address, (2-12), (0), (14), or n
you may specify either the address of a field containing the hexadecimal value of the record's key length, or you may specify the record's key length using the low-order byte of a register or immediate data (n). Specify a decimal digit for n (immediate data).

DD=addr—RX-type address, (2-12), (0), (14), or n
you may specify either the address of a field containing the hexadecimal value of the record's data length, or you may specify the record's data length using the low-order two bytes of a register or immediate data (n). Specify a decimal digit for n (immediate data).

REGSAVE={YES|NO}

YES

specifies registers 1 through 14 are saved and restored in the caller-provided save area (pointed to by register 13) across the TRKCALC call. Otherwise, registers 1, 9, 10, 11, and 14 are modified. Registers 0 and 15 are always modified by a TRKCALC call.

NO

specifies registers are not saved across a TRKCALC call. NO is the default.

MF=I

specifies storage definition for the TRKCALC parameter list and parameter list initialization, using the given keywords, then calling the TRKCALC function. MF=I is the default.

TRKCALC—EXECUTE FORM

A remote parameter list is referred to and can be modified by the execute form of the TRKCALC macro. The TRKCALC routine is called. The description of the standard form of the macro provides the explanation of the function of each operand.

[<u>symbol</u>]	TRKCALC	[FUNCTN={TRKBAL TRKCAP}] [{,DEVTAB={ <u>addr</u> *}] ,UCB={ <u>addr</u> *},TYPE={ <u>addr</u> *}] [,BALANCE={ <u>addr</u> *}] [,REMOVE={YES NO}] [,MAXSIZE={YES NO}] [{,RKDD= <u>addr</u> },R= <u>addr</u> ,K= <u>addr</u> ,DD= <u>addr</u> }] [,REGSAVE={YES NO}] ,MF=(E, <u>addr</u>)
-------------------	---------	--

FUNCTN={TRKBAL|TRKCAP}

is coded as shown in the standard form. If this keyword is omitted, any specification of REMOVE, MAXSIZE, LAST, and the RX form of BALANCE is ignored. In addition, DEVTAB is assumed if UCB is coded and a failure occurs, or if TYPE is specified. When you use FUNCTN, one of the keywords (DEVTAB, UCB, or TYPE) must be specified to provide an information source.

DEVTAB=addr|*—RX-type address, (2-12), (0), (14)

is coded as shown in the standard form except for the * subparameter. Specify an * when you have inserted the address of the device characteristics table entry (DCTE) in the parameter list.

UCB=addr|*—RX-type address, (2-12), (0), (14)

is coded as shown in the standard form except for the * subparameter. Specify an * when you have inserted the address of the UCB in the parameter list.

TYPE=addr|*—RX-type address, (2-12), (0), (14)

is coded as shown in the standard form except for the * subparameter. Specify an * when you have inserted the address of the UCB type (UCBTYP) in the parameter list.

BALANCE=addr|*—RX-type address, (2-12), (0), (14)

is coded as shown in the standard form except for the * subparameter. Specify an * when you have inserted the balance in the parameter list.

REMOVE={YES|NO}

is coded as shown in the standard form.

MAXSIZE={YES|NO}

is coded as shown in the standard form.

RKDD=addr—RX-type address, (2-12), (0), (14)

is coded as shown in the standard form.

R=addr—RX-type address, (2-12), (0), (14) or n

is coded as shown in the standard form.

K=addr—RX-type address, (2-12), (0), (14), or n

is coded as shown in the standard form.

DD=addr—RX-type address, (2-12), (0), (14), or n

is coded as shown in the standard form.

REGSAVE={YES|NO}

is coded as shown in the standard form.

MF=(E,addr)

This operand specifies that the execute form of the TRKCALC macro instruction and an existing data management parameter list are used.

E
Coded as shown.

addr—RX-type address, (0), (1), (2-12), or (14)
specifies an in-storage address of the parameter list.

TRKCALC—LIST FORM

The list form of the TRKCALC macro constructs an empty, in-line parameter list. By coding only MF=L, you construct a parameter list, and the actual values can be supplied by the execute form of the TRKCALC macro. Any parameters other than MF=L are ignored.

[symbol]	TRKCALC	MF=L
----------	---------	------

TRKCALC—DSECT ONLY

This call gives a symbolic expansion of the parameter list for the TRKCALC macro. No DSECT statement is generated. If a name is specified on the macro call, it applies, after any necessary boundary alignment, to the beginning of the list. The macro-generated symbols all begin with "STAR".

[symbol]	TRKCALC	MF=D
----------	---------	------

INPUT REGISTER USAGE FOR ALL FORMS OF 'MF'

Registers 0, 2 through 12, and 14 are available to provide input for keywords.

Register 1 is used only to provide the address of the parameter list for an MF=E call.

Register 13 may be used as input for keywords, if REGSAVE=YES is not specified.

Register 15 is used as a work register to build the TRKCALC parameter list for the MF=E call; it is not available as an input register.

OUTPUT FROM TRKCALC

FUNCTN=TRKBAL

Register 15=0

The record fits on the track. Register 0 and STARBAL contain the new track balance.

Register 15=4

Record does not fit on the track. If MAXSIZE=YES is specified, a partial record does not fit either. Register 0 and STARBAL are set to zero.

Register 15=8

Record does not fit on the track. MAXSIZE=YES is specified, and a partial record does fit. Register 0 and STARBAL are set to the maximum number of data bytes that fit on the remainder of the track with the specified keylength.

Note: The keylength is excluded from the count of maximum data bytes.

STARBAL

This is the track balance field of the TRKCALC parameter list. This field is first set to the track capacity if R=1, or to the supplied BALANCE value if R≠1, or to the calculated balance if R≠1 and BALANCE are omitted. STARBAL is updated to the new track balance if the record fits; otherwise, STARBAL is left with the input track balance value.

FUNCTN=TRKCAP

Register 15=0

Register 0 contains the number of records that fit on the track if R = 1, or the number of records that fit on the remainder of the track if R ≠ 1.

Register 15=4

No records of the length specified fit on a full track (R = 1) or a partial track (R ≠ 1). Register 0 is set to zero.

STARBAL

This is the track balance field of the TRKCALC parameter list. This field is first set to the track capacity if R=1, or to the supplied BALANCE value if R≠1, or to the calculated balance if R≠1 and BALANCE are omitted.

RETURN CODES FROM TRKCALC

The TRKCALC macro passes a return code in register 15. The return codes and their meanings are as follows:

Contents Meaning

- 00(X'00') Indicates that register 0 contains the new track balance
- 04(X'04') Indicates that the record did not fit (register 0 = 0)
- 08(X'08') Indicates that the record did not fit (register 0 contains the maximum data length that does fit)

TRKCALC MACRO EXAMPLES

In this example, TRKCALC is coded to determine how many records of a given size with 10-byte keys fit on a 3380 track. After issuing the macro, the number of records is saved in NUMREC:

```
TRKCALC  FUNCTN=TRKCAP,TYPE=UTYPE,R=1,K=10,DD=DL,MF=(E,(1))
.
.
.
          ST    0,NUMREC    SAVE NUMBER OF RECORDS
.
.
          DL    DC    H'xxxx'    DATA LENGTH
          UTYPE DC    X'0E'
          NUMREC DS    F          MAX # OF RECORDS
```

In this example, TRKCALC is coded to determine whether another record can fit on a track of a 3380, given a track balance.

```
TRKCALC  FUNCTN=TRKBAL,TYPE=UTYPE,R=REC,K=KL,DD=DD,BALANCE=BAL,
          MAXSIZE=YES,MF=(E,(1))
.
.
          UTYPE DC    X'0E'
          REC   DC    X'xx'
          KL    DC    X'xx'
          DD    DC    H'xxxx'
          BAL   DC    H'xxxx'
```

ALLOCATING A DASD DATA SET

The REALLOC macro builds a parameter list and issues SVC 32 to allocate a new data set based on a partial DSCB that describes the attributes of that data set. You can use the OBTAIN macro to get the format-1 DSCB of another data set and use it as a model for the new data set's DSCB which REALLOC will construct and write in the VTOC.

The maximum number of extents that may be allocated are determined by the type of data set requested as defined by the data set organization (DS1DSORG) bytes and the data set indicator (DS1DSIND) byte in the partial DSCB. If the DS1DSORG indicates a VSAM data set organization and DS1DSIND indicates the data set is cataloged in an integrated catalog facility (ICF) catalog, the maximum number of extents is 123. Otherwise, the maximum number of extents is 16.

Note: User label data sets, ISAM data sets, and absolute track allocated data sets are not supported by the REALLOC macro. If a VSAM data set or data space is requested, REALLOC does not interface with VSAM or integrated catalog facility catalog management.

The DS1SCALO field of the partial format-1 DSCB has a high order flag byte that describes the type of request and a 3-byte field containing the secondary allocation quantity. The following describes the flag byte:

Contents Meaning

X'C0'	Cylinder request
X'C8'	Cylinder with CONTIG request
X'80'	Track request
X'88'	Track with CONTIG request
X'40'	Average block length request
X'41'	Average block length with ROUND request
X'48'	Average block length with CONTIG request
X'49'	Average block length with CONTIG and ROUND request

Any settings other than the above are ignored.

The REALLOC macro may be coded in the execute, DSECT, and list forms, but not the standard form. The calling program may be in supervisor or problem program state, and may be running in any key. The calling program must be APF authorized.

REALLOC—EXECUTE FORM

The format of the REALLOC macro in execute form is:

[symbol]	REALLOC	MF=(E,addr) ,DSSIZE=addr (reg) ,PDSCB=addr ,UCB=addr [,MINAU=addr (reg) [,PDSDIR=addr (reg)
----------	---------	--

MF=(E,addr)
specifies that the execute form of the macro and an existing REALLOC parameter list are used.

addr—RX-type address, (0-12)
specifies an in-storage address of the REALLOC parameter list.

E
Code as shown.

DSSIZE=addr|(reg)
specifies the size of the data set to be allocated in tracks. If a cylinder request (X'C0' in the flag byte of DS1SCAL0) or average block with round request (X'41') is made, the number of tracks specified is rounded up to the next full cylinder, if necessary.

You may not specify the DSSIZE in terms of average block size, even though the original data set may have been allocated with the number of average blocks (X'40').

addr—RX-type address
specifies an in-storage address of a full word containing the data set size.

(reg)
specifies a register containing the size of the data set. Valid registers are 0 and 2-12.

PDSCB=addr—RX-type address, (0), (2-12)
specifies the address of the partial DSCB. The partial DSCB is comprised of the first 98 bytes of a format-1 DSCB. The first 44 bytes contains the data set name to be allocated. The contents of the partial DSCB are used, unchanged, in constructing the format-1 DSCB. Only the field DS1NOEPV (number of extents on the volume) of the partial format-1 DSCB is modified by allocation to reflect the actual number of extents allocated.

UCB=addr—RX-type address, (0), (2-12)
specifies the address of the UCB of the volume where the data set is to be allocated. The volume must be mounted and you must ensure that it remains mounted.

MINAU=addr|(reg)
specifies the size of the minimum allocation unit in tracks. All primary extents for this data set are in multiples of this minimum allocation unit. This value does not apply to subsequent extensions of the data set. If the partial DSCB indicates the data set is to be allocated in cylinders (X'C0' or average block with round request (X'41')), this parameter is ignored.

addr—RX-type address
specifies an in-storage address of a full word containing the minimum allocation unit.

(reg)
specifies a register containing the minimum allocation unit. Valid registers are 0 and 2-12.

PDSDIR=addr|(reg)
specifies the number of 256 byte directory blocks for a partitioned data set (PDS). This is a required keyword if the DS1DSORG indicates a partitioned data set. Otherwise, it is ignored.

addr—RX-type address
specifies an in-storage address of a full word containing the number of 256 byte PDS directory blocks.

(reg)
specifies a register containing the number of 256 byte PDS directory blocks. Valid registers are 0 and 2-12.

REALLOC—DSECT ONLY

The DSECT form of REALLOC is specified as follows:

[symbol]	REALLOC	MF=D
----------	---------	------

An example of the DSECT form expansion is:

REALPL	REALLOC MF=D	
REALPL	DSECT	DSECT FOR PARAMETER LIST
RALPLID	DS CL4	EBCDIC 'REAL' FOR REALLOC
RALNGTH	DS AL2	LENGTH OF PARAMETER LIST
RAERRCDE	DS H	ERROR CODE RETURNED FROM
*		ALLOCATE (SVC 32)
RALRSVD	DS F	RESERVED
RALDSSZ	DS F	DATA SET SIZE
RALMAU	DS F	MINIMUM ALLOCATION UNIT
RALPDSCB	DS A	PARTIAL DSCB POINTER
RALUCB	DS A	UCB POINTER
RALDQTY	DS F	PDS DIRECTORY QUANTITY
RALEND	EQU *	END OF PARAMETER LIST
RALENGTH	EQU RALEND-REALPL	LENGTH OF PARAMETER LIST

REALLOC—LIST FORM

The list form of the REALLOC macro is specified as follows:

[symbol]	REALLOC	MF=L ,DSSIZE=addr (reg) ,PDSCB=addr ,UCB=addr [,MINAU=addr (reg) [,PDSDIR=addr (reg)
----------	---------	---

Refer to the execute form for an explanation of the parameters.

An example of the list form expansion is:

REALPL	REALLOC MF=L	
REALPL	CNOP 0,4	
REALPL	EQU *	
	DC CL4'REAL'	EBCDIC 'REAL' FOR REALLOC
	DC AL2(32)	LENGTH OF PARAMETER LIST
	DC H'0'	ERROR CODE RETURNED FROM
*		ALLOCATE (SVC 32)
	DC F'0'	RESERVED
	DC F'0'	DATA SET SIZE
	DC F'0'	MINIMUM ALLOCATION UNIT
	DC A(0)	PARTIAL DSCB POINTER
	DC A(0)	UCB POINTER
	DC F'0'	PDS DIRECTORY QUANTITY
RAL01E	EQU *	END OF PARAMETER LIST

RETURN CODES FROM REALLOC

Control returns to the instruction following the SVC 32 generated by the REALLOC macro. If the data set was successfully allocated, register 15 contains zeros. Otherwise, register 15 contains one of the following return codes²:

² This is a cumulative list of DADSM allocation return codes. Some of these codes may not apply to the REALLOC macro.

Code	Meaning
004(X'04')	Data set name of request already exists on this volume. Initial allocation not possible under the name given.
008(X'08')	No room available in the VTOC or VTOC index.
012(X'0C')	One of the following errors was encountered: <ul style="list-style-type: none"> • Permanent I/O error • Error returned by CVAF
020(X'14')	Requested quantity not available.
028(X'1C')	ISAM DSORG is not supported.
048(X'30')	Register 0 contains a reason code indicating one of the following errors: <ul style="list-style-type: none"> • Reason Code 1 - Invalid REALLOC parmlist ID. • Reason Code 2 - Invalid REALLOC parmlist length.
052(X'34')	Invalid partial DSCB pointer.
056(X'38')	Not enough space on volume for directory.
072(X'48')	DOS VTOC cannot be converted to an OS VTOC.
116(X'74')	User labels not supported.
120(X'78')	DSSIZE=0 and MINAU is greater than 0.
124(X'7C')	DSSIZE is not a multiple of MINAU.
128(X'80')	Directory space requested is larger than primary space.
148(X'94')	Overlapping extents in the VTOC.
152(X'98')	Overlapping DOS split cylinder extents in the VTOC.
156(X'9C')	DADSM allocation terminated due to possible VTOC errors.
164(X'A4')	Allocation terminated due to DOS stacked pack format.
168(X'A8')	RACF DEFINE failed, data set already defined.
172(X'AC')	User not authorized to RACF define data set.
176(X'B0')	Installation exit rejected this request with a return code of 8.
180(X'B4')	Installation exit rejected this request with a return code of 4.

MESSAGE DISPLAYS ON THE IBM 3480 MAGNETIC TAPE SUBSYSTEM

The MSGDISP macro displays a message on the IBM 3480. With MSGDISP, you can specify the message to be displayed and how to display it (for example, steady or flashing). The six main parameters of the macro and their functions are:

MOUNT Displays an "M" in position 1 of the display area during a mount request until a volume is loaded and made ready. The "M" is followed by the volume serial number and label type.

- VERIFY** Shows that a volume has been accepted by displaying its serial number and label type in positions 2 through 8.
- RDY** Displays text in positions 2 through 7 while a data set is open.
- DEMOUNT** Displays a volume disposition indicator in position 1 until a volume is demounted.
- RESET** Clears the display area.
- GEN** Provides the full range of display options, including the option to alternate two messages.

All except the RDY parameter require that you be in supervisor state, have a storage protect key of 0 through 7, or be authorized by the authorized program facility.

The MSGDISP macro generates a parameter list as input to an SVC routine.

MSGDISP may be coded in the standard, execute, and list forms.

The formats for specifying MSGDISP with the six main parameters, and the return codes generated by MSGDISP, are given in the sections that follow.

MSGDISP—DISPLAYING A MOUNT MESSAGE

The format for specifying MSGDISP with the MOUNT parameter is:

[symbol]	MSGDISP	MOUNT ,UCB=(reg) [,LABEL={'A' 'N' 'S' 'X' addr}] [,MF={L (E,addr)}] [,SER={'volser' addr}] [,TEST={NO YES}] [,WAIT={NO YES}]
----------	---------	--

MOUNT
displays an "M" in position 1 of the display area during a mount request. The "M" is followed by a volume serial number and label type. The display flashes on and off until a volume is loaded and ready. If the device is ready at the time a mount request is issued, the "M" is not displayed.

UCB=(reg)—(2-12)
specifies a register containing the UCB address for the device.

LABEL={'A'|'N'|'S'|'X'|addr}
displays the label type of the mounted volume in position 8. If you specify an unknown label type other than a blank, a "?" is displayed.

'A'
specifies ISO/ANSI/FIPS (AL) or ISO/ANSI/FIPS with user labels (AUL). Specify in apostrophes.

'N'
specifies no labels (NL), LTM (DOS), or bypass label processing (BLP). Specify in apostrophes.

'S'
specifies IBM Standard (SL) or IBM Standard with user labels (SUL). Specify in apostrophes.

'X'
specifies nonstandard labels (NSL). Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)
specifies an in-storage address of an area containing an "A", "N", "S", or "X" (see the following explanations of these characters). For MF=L, you may only specify an A-type address.

MF={L|(E,addr)}
specifies either the execute or the list form of MSGDISP. If you do not specify this parameter, the standard form of the macro is used.

L
specifies the list form of MSGDISP. This generates a parameter list that can be used as input to the execute form. The execute form can modify the parameter list.

(E,addr)
specifies that the execute form of the macro and an existing parameter list are used.

addr—RX-type address, (1), or (2-12)
specifies an in-storage address of the parameter list.

SER={'volser'|addr}
specifies the serial number of the volume to be mounted. The serial number is displayed in positions 2 through 7. If you do not specify SER, the system supplies the volume serial number. If the serial number is not available, a scratch volume is used, unless the volume use attribute indicates a default of "PRIVAT".

'volser'
specifies the volume serial number as a literal. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)
specifies an in-storage address of the volume serial number. For MF=L, you may only specify an A-type address.

TEST={NO|YES}
specifies whether to test the UCB to determine if the device is capable of displaying messages.

NO
specifies that the SVC routine will test the UCB.

YES
specifies testing the UCB before the SVC call.

Note: TEST=YES requires you to include the UCB mapping macro (IEFUCBOB) in the source code.

WAIT={NO|YES}
specifies when control is returned to you.

NO
specifies that control is to be returned before I/O is complete. I/O return codes are not returned, and I/O errors are recorded in the same manner as any permanent error by the error recovery procedure.

YES
specifies that control is to be returned after I/O is complete.

MSGDISP—DISPLAYING A VERIFY MESSAGE

The format for specifying MSGDISP with the VERIFY parameter is:

[symbol]	MSGDISP	VERIFY ,UCB=(reg) [,LABEL={'A' 'N' 'S' 'X' addr}] [,MF={L (E,addr)}] [,SER={'volser' addr}] [,TEST={NO YES}] [,WAIT={NO YES}]
----------	---------	---

VERIFY

displays the serial number and label type of a volume that has been accepted in positions 2 through 8. Position 1 remains blank. The display lasts until the next display request is executed.

UCB=(reg)—(2-12)

specifies a register containing the UCB address for the device.

LABEL={'A'|'N'|'S'|'X'|addr}

specifies label type of the mounted volume in position 8 of the display. If an unknown label type other than a blank is specified, a "?" is displayed.

'A'

specifies ISO/ANSI/FIPS (AL) or ISO/ANSI/FIPS with user (AUL) labels. Specify in apostrophes.

'N'

specifies no labels (NL), LTM (DOS), or bypass label processing (BLP). Specify in apostrophes.

'S'

specifies IBM Standard (SL) or IBM Standard with user (SUL) labels. Specify in apostrophes.

'X'

specifies nonstandard (NSL) labels. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)

specifies an in-storage address of an area containing an "A", "N", "S", or "X" (see explanations below for these characters). For MF=L, you may only specify an A-type address.

MF={L|(E,addr)}

specifies either the execute or list form of MSGDISP. If you do not specify this parameter, the standard form of the macro is used.

L

specifies the list form of MSGDISP. This generates a parameter list that can be used as input to the execute form. The execute form can modify the parameter list.

(E,addr)

specifies that the execute form of the macro and an existing parameter list is to be used.

addr—RX-type address, (1), or (2-12)

specifies an in-storage address of the parameter list.

SER={'volser'|addr}

specifies the serial number of the volume that has been verified. The serial number displays in positions 2 through 7. If you do not specify SER, the system supplies

the volume serial number. If the serial number is not available, a scratch volume is used, unless the volume use attribute indicates a default of "PRIVAT".

'volser'
specifies the volume serial number as a literal. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12) specifies an in-storage address of the volume serial number. For MF=L, you may only specify an A-type address.

TEST={NO|YES}
specifies whether to test the UCB to determine if the device is capable of displaying messages.

NO
specifies that the SVC routine will test the UCB.

YES
specifies testing the UCB before the SVC call.

Note: TEST=YES requires you to include the UCB mapping macro (IEFUCBOB) in the source code.

WAIT={NO|YES}
specifies when control is to be returned to you.

NO
specifies that control is to be returned before I/O is complete. I/O return codes are not returned, and I/O errors are recorded in the same manner as any permanent error by the error recovery procedure.

YES
specifies that control is to be returned after I/O is complete.

MSGDISP—DISPLAYING A READY MESSAGE

The format for specifying MSGDISP with the RDY parameter is:

[symbol]	MSGDISP	RDY ,DCB=addr [,MF={L (E,addr)}] [,TXT={'msgtxt' addr}]
-----------------	----------------	--

RDY
displays the text supplied in the TXT parameter in positions 2 through 7 while the data set is open. The display is steady (not flashing) and is enclosed in parentheses. The display is also written to the tape pool console (routing code 3, descriptor code 7).

DCB=addr
specifies the address of a DCB opened to a data set on the mounted volume. If multiple devices are allocated, the message display is directed to the one containing the volume currently in use.

Note: If multiple devices or multiple volumes are allocated, you may update a message display after an end-of-volume condition by using the EOVS exit specified in a DCB exit list. In the case of a concatenated data set with unlike characteristics, the DCB OPEN exit may be used to update the display.

addr—RX-type address, A-type address, or (2-12) specifies an in-storage address of the opened DCB. For MF=L, you may only specify an A-type address.

MF={L|(E,addr)}

specifies either the execute or list form of MSGDISP. If this parameter is not specified, the standard form of the macro is used.

L

specifies the list form of MSGDISP. This generates a parameter list that can be used as input to the execute form. The execute form can modify the parameter list.

(E,addr)

specifies that the execute form of the macro and an existing parameter list is to be used.

addr—RX-type address, (1), or (2-12)

specifies an in-storage address of the parameter list.

TXT={'msgtxt'|addr}

specifies up to six characters to display in positions 2 through 7 of the display. If you do not specify TXT, blanks are displayed.

'msgtxt'

specifies the text as a literal. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)

specifies an in-storage address of an area containing the text to be displayed. For MF=L, you may only specify an A-type address.

MSGDISP—DISPLAYING A DEMOUNT MESSAGE

The format for specifying MSGDISP with the DEMOUNT parameter is:

[symbol]	MSGDISP	DEMOUNT
		,UCB=(reg) [,DISP={'D' 'K' 'R' addr}] [,MF={L (E,addr)}] [,MLABEL={'A' 'N' 'S' 'X' addr}] [,MSER={'volser-to-mount' addr}] [,SER={'volser' addr}] [,TEST={NO YES}] [,WAIT={NO YES}]

DEMOUNT

Displays a volume disposition indicator in position 1 until the volume is demounted. Optionally, you may display the serial number of the volume to be demounted at the same time. The display flashes on and off. If a volume is not mounted on the device when the display request is executed, blanks are displayed.

The demount message may be displayed alternately (flashing) with a mount message for the next volume by specifying the MSER parameter.

UCB=(reg)—(2-12)

specifies a register containing the UCB address for the device.

DISP={'D'|'K'|'R'|addr}

specifies the character to display in position 1, representing the volume disposition.

'D'

Demount a public volume. Specify in apostrophes.

Note: "D" also displays when you specify an invalid character or when the volume use attribute is unknown (as in an automatic volume recognition (AVR) error when reading a label).

'K'
Keep a private volume and return it to the library. Specify in apostrophes.

'R'
Retain a private volume near the device for further use. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)
specifies an in-storage address of an area containing a "D", "K", or "R". For MF=L, you may only specify an A-type address.

MF={L|(E,addr)}
specifies either the execute or list form of MSGDISP. If you do not specify this parameter, the standard form of the macro is used.

L
specifies the list form of MSGDISP. This generates a parameter list that can be used as input to the execute form. The execute form can modify the parameter list.

(E,addr)
specifies that the execute form of the macro and an existing parameter list is to be used.

addr—RX-type address, (1), or (2-12)
specifies an in-storage address of the parameter list.

MLABEL={'A'|'N'|'S'|'X'|addr}
displays the label type of the volume to be loaded and made ready following a demount, in position 8. If you specify an unknown label type other than a blank, a "?" is displayed. You may only specify this parameter if you also specify the MSER parameter.

'A'
specifies ISO/ANSI/FIPS (AL) or ISO/ANSI/FIPS with user (AUL) labels. Specify in apostrophes.

'N'
specifies no labels (NL), LTM (DOS), or bypass label processing (BLP). Specify in apostrophes.

'S'
specifies IBM Standard (SL) or IBM Standard with user (SUL) labels. Specify in apostrophes.

'X'
specifies nonstandard (NSL) labels. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)
specifies an in-storage address of an area containing an "A", "N", "S", or "X" (see the following explanations of these characters). For MF=L, you may only specify an A-type address.

MSER={volser-to-mount'|addr}
displays the mount message for the next volume alternately (flashing) with the demount message. The display continues until you demount the current volume. At that time, the mount message will display (flashing) until you load the volume and make the device ready. If no volume is mounted at the time the demount and mount messages are executed,

only the mount message will display (flashing) until the volume is loaded and ready.

'volser-to-mount'

specifies the volume serial number, as a literal, of the volume to be mounted. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)
specifies an in-storage address of the volume serial number of the volume to be mounted. For MF=L, you may only specify an A-type address.

SER={volser'|addr}

specifies the serial number of the volume to be demounted. The serial number is displayed in positions 2 through 7. If you do not specify SER, the system supplies the volume serial number. If the serial number is not available, a scratch volume is used, unless the volume use attribute indicates a default of "PRIVAT".

'volser'

specifies the volume serial number as a literal. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)
specifies an in-storage address of the volume serial number. This parameter is not valid for the MF=L form. For MF=L, you may only specify an A-type address.

TEST={NO|YES}

specifies whether to test the UCB to determine if the device is capable of displaying messages.

NO

specifies that the SVC routine will test the UCB.

YES

specifies testing the UCB before the SVC call.

Note: TEST=YES requires you to include the UCB mapping macro (IEFUCBOB) in the source code.

WAIT={NO|YES}

specifies when control is to be returned to you.

NO

specifies that control is to be returned before I/O is complete. I/O return codes are not returned, and I/O errors are recorded in the same manner as any permanent error by the error recovery procedure.

YES

specifies that control is to be returned after I/O is complete.

MSGDISP—RESETTING THE MESSAGE DISPLAY

The format for specifying MSGDISP with the RESET parameter is:

[symbol]	MSGDISP	RESET ,{UCB=(reg),UCBL=addr} [,MF={L (E,addr)}] [,TEST={NO YES}] [,WAIT={NO YES}]
----------	---------	---

RESET

clears all existing data on the display. If you specify WAIT=NO and the last service requested was a demount, the display is not cleared.

After being cleared, the display will show the device's internal status message (for example, a message indicating that the device is ready).

UCB=(reg)—(2-12)

specifies a register containing the UCB address for the device.

UCBL=addr—RX-type address, A-type address, (0), or (2-12) specifies the address of the list containing a maximum of 64 words. Each word in the list contains the address of a UCB representing a device whose display is to be reset. The end of the list is indicated by a '1' in the high-order bit of the last address in the list. If an error is encountered while processing the list, register 1 points to the associated UCB when you regain control.

You cannot specify UCBL with TEST=YES and WAIT=NO.

MF={L|(E,addr)}

specifies either the execute or the list form of MSGDISP. If you do not specify this parameter, the standard form of the macro is used.

L

specifies the list form of MSGDISP. This generates a parameter list that can be used as input to the execute form. The execute form can modify the parameter list.

(E,addr)

specifies that the execute form of the macro and an existing parameter list is to be used.

addr—RX-type address, (1), or (2-12)

specifies the address of the parameter list.

TEST={NO|YES}

specifies whether to test the UCB to determine if the device is capable of displaying messages.

NO

specifies that the SVC routine will test the UCB.

YES

specifies testing the UCB before the SVC call. You cannot specify TEST=YES if you also specify the UCBL parameter.

Note: TEST=YES requires you to include the UCB mapping macro (IEFUCBOB) in the source code.

WAIT={NO|YES}

specifies when control is to be returned to you.

NO

specifies that control is to be returned before I/O is complete. I/O return codes are not returned, and I/O errors are recorded in the same manner as any permanent error by the error recovery procedure.

You cannot specify WAIT=NO if you also specify the UCBL parameter.

YES

specifies that control is to be returned after I/O is complete.

Note: Demount messages can be reset only if WAIT=YES is specified.

MSGDISP—PROVIDING THE FULL RANGE OF DISPLAY OPTIONS

The format for specifying MSGDISP with the GEN parameter is:

[symbol]	MSGDISP	GEN ,UCB=(reg) [,FLASH={STEADY STEADY2 BLINK BLINK2 ALT}] [,MF={L (E,addr)}] [,TEST={NO YES}] [,TXT={'msgtxt' addr}] [,TXT2={'altmsgtxt' addr}] [,VOL={STATIC REMOVE INSERT SWAP}] [,WAIT={NO YES}]
----------	---------	---

GEN

specifies the full range of display options.

UCB=(reg)—(2-12)

specifies a register containing the UCB address for the device.

FLASH={STEADY|STEADY2|BLINK|BLINK2|ALT}

specifies message display mode.

Note: If you specify VOL=SWAP, messages will always be displayed as if you had specified FLASH=ALT.

STEADY

specifies that the primary message (TXT) is to be displayed without flashing.

STEADY2

specifies that the alternate message (TXT2) is to be displayed without flashing.

BLINK

specifies that the primary message (TXT) flashes on and off at a rate of approximately two seconds on and one-half second off.

BLINK2

specifies that the alternate message (TXT2) flashes on and off at a rate of approximately two seconds on and one-half second off.

ALT

specifies that the primary and alternate messages (TXT and TXT2) flash on and off alternately at a rate of approximately two seconds on and one-half second off.

MF={L|(E,addr)}

specifies either the execute or the list form of MSGDISP. If you do not specify this parameter, the standard form of the macro is used.

L
specifies the list form of MSGDISP. This generates a parameter list that can be used as input to the execute form. The execute form can modify the parameter list.

(E,addr)
specifies that the execute form of the macro and an existing parameter list is to be used.

addr
specifies an in-storage address of the parameter list. Specify either an RX-type address or a register in the range of 2 through 12.

TEST={NO|YES}
specifies whether to test the UCB to determine if the device is capable of displaying messages.

NO
specifies that the SVC routine will test the UCB.

YES
specifies testing the UCB before the SVC call.

Note: TEST=YES requires you to include the UCB mapping macro (IEFUCBOB) in the source code.

TXT={'msgtxt'|addr}
specifies 8 characters to be shown in positions 1 through 8 of the display. If you do not specify TXT, blanks are displayed.

'msgtxt'
specifies the 8 characters as literals. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)
specifies an in-storage address of an area containing the 8 characters. For MF=L, you may only specify an A-type address.

TXT2={'altmsgtxt'|addr}
specifies 8 alternate characters to display in positions 1 through 8 of the display. If you do not specify TXT2, blanks are displayed.

'alt-msgtxt'
specifies the 8 characters as literals. Specify in apostrophes.

addr—RX-type address, A-type address, or (2-12)
specifies an in-storage address of an area containing the 8 characters. For MF=L, you may only specify an A-type address.

VOL={STATIC|REMOVE|INSERT|SWAP}
specifies message display mode, based on volume status.

STATIC
specifies that messages will display without regard to volume status until the next message request is executed, or until the next command initiates volume movement.

REMOVE
specifies that messages will display until the current volume is demounted. This parameter is ignored if a volume is not mounted when the request is executed.

INSERT
specifies that messages will display until a volume is present, the tape threaded, and the active/inactive switch is in the active position. This parameter is

ignored if a volume is loaded and ready when the request is executed.

SWAP

specifies that messages will always display as if FLASH=ALT were specified. The data from TXT and TXT2 displays alternately (flashing) until the current volume has been demounted. Then only TXT2 will display (flashing) until a new volume is loaded and ready. If no volume is mounted when this parameter is specified, only TXT2 data will display (flashing) until a new volume is loaded and ready.

WAIT={NO|YES}

specifies when control is to be returned to you.

NO

specifies that control is to be returned before I/O is complete. I/O return codes are not returned, and I/O errors are recorded in the same manner as any permanent error by the error recovery procedure.

YES

specifies that control is to be returned after I/O is complete.

RETURN CODES FROM MSGDISP

When the system returns control to the problem program, the low-order byte of register 15 contains a return code. The low-order byte of register 0 may contain a reason code as follows:

Return Code (15)	Reason Code (0)	Meaning
00(X'00')		Successful completion.
04(X'04')		Device does not support MSGDISP.
08(X'08')		Unauthorized request (failed TESTAUTH for proper authority level) or invalid input parameters (including DCB or UCB).
08(X'08')	01(X'01')	Invalid parameter.
	02(X'02')	Invalid DCB or DEBCHK error.
	03(X'03')	Environmental error.
	04(X'04')	Authorization violation.
	05(X'05')	Invalid UCB.
	06(X'06')	Invalid request.
	11(X'0B')	Unsuccessful ESTAE macro call.
	12(X'0C')	Unsuccessful GETMAIN request.
12(X'0C')		I/O error (I/O Supervisor posted the request for an error).

Note: An I/O error occurs for load display if the drive display has a hardware failure.

If you get return code X'04' or X'0C' on a RESET UCBL operation, register 1 points to the UCB associated with the error when you regain control.

CHAPTER 7. MAINTAINING SYS1.IMAGELIB

This chapter describes how to maintain the system image library (SYS1.IMAGELIB) UCS and FCB images for the IBM 1403, 3203, and 3211 Printers. It also describes how to maintain FCB images for the IBM 4245 Printer, the UCS image table in SYS1.IMAGELIB for the 3262 Model 5, 4245, and 4248 Printers, and how to retrieve an FCB image from SYS1.IMAGELIB for modification.

The IEBIMAGE utility program is used to create and maintain control modules for the IBM 3800 Printing Subsystem: character arrangement table modules, graphic character modification modules, copy modification modules, library character set modules, and FCB modules.

IEBIMAGE can also be used to create and maintain FCB modules for the 4248 Printer. FCB modules created for the 4248 can also be used with the 3262 Model 5 Printer; however, the 3262 Model 5 does not support variable printer speeds or the horizontal copy feature of the 4248. For more information about IEBIMAGE, see Utilities.

To use the information presented in this chapter, you should be familiar with the subjects of the following publications:

- Data Administration: Macro Instruction Reference describes the SETPRT macro that can specify the UCS and/or FCB images to be used.
- JCL describes the UCB and FCB parameters of the DD statement that are processed at OPEN time.
- IBM 2821 Control Unit Component Description contains information on creating a user-designed chain/train for the 1403 Printer.
- IBM 3203 Printer Component Description and Operator's Guide contains information on creating a user-designed train for the 3203 Printer.
- IBM 3211 Printer, 3216 Interchangeable Train Cartridge, and 3811 Printer Control Unit Component Description and Operator's Guide contains information on creating a user-designed train for the 3211 Printer.
- System Programming Library: JES2 Initialization and Tuning or Network Job Entry Facility for JES2 contains reference information for JES2.
- JES3 Initialization and Tuning contains reference information for JES3.
- The IBM 3262 Model 5, 4245 and 4248 Printers have UCS images supplied in SYS1.IMAGELIB. The IBM-supplied image tables are shown in Figure 29 on page 159 and Figure 32 on page 163. For a list of all UCS images available, see:
 - IBM 3262 Printer Model 5 Product Description, containing information on band IDs for the 3262 Model 5 Printer.
 - IBM 4245 Printer Model 1 Component Description and Operator's Guide, containing information on band IDs for the 4245 Printer.
 - IBM 4248 Printer Description, containing information on band IDs for the 4248 Printer.

The SPZAP service aid can be used to display and modify an existing member of SYS1.IMAGELIB. Use of SPZAP on load modules is described in Service Aids.

UCS IMAGES IN SYS1.IMAGELIB

Most IBM standard character set images are included in SYS1.IMAGELIB at system generation time, through the DATAMGT macro and an IODEVICE macro for the specified printer. (For details on the DATAMGT and IODEVICE macros, see System Generation Reference.) The standard character set images for the 1403, 3203, and 3211 Printers are shown below.

Printer	Images
1403 or 3203	AN, HN, PCAN, PCHN, PN, QNC, QN, RN, SN, TN, XN, YN
3211	A11, G11, H11, P11, T11

For the 3262 Model 5, 4245, and 4248 Printers, no UCS images are supplied in SYS1.IMAGELIB. Instead, a new UCS image is loaded into the buffer when the machine is powered on or the operator mounts a new band. See "Adding a UCS Image Name/Alias to a UCS Image Table" on page 162 for information on how to access UCS images that are not supplied in SYS1.IMAGELIB.

The 3262 Model 5, 4245 and 4248 Printers also load a default FCB image when the machine is powered on. The default FCB for the 4248 is the last FCB loaded into the buffer. For the 4245, the default FCB is an 11-inch form with 6 LPI and a Channel 1 on the first print line. For the 3262 Model 5, the default FCB is an 11-inch form with 6 LPI, a Channel 1 on the third print line, and a Channel 12 on line 64.

The alias names are defined for most installation-standard print chains/trains/bands installable on a given printer. Alias names are included in SYS1.IMAGELIB (in the UCS image table) at system generation time, with the real name of each image.

Some print chains/trains/bands, such as SN and G11, do not have alias names because there is no equivalent chain/train/band on other printers. You can assign an alias for these chains/trains/bands with the ALIAS statement of the linkage editor. (For more information on the ALIAS statement, see Linkage Editor and Loader.) For the 3262 Model 5, 4245 or 4248 Printer, you can also add an alias name by modifying an entry in the UCS image table. See "Adding a UCS Image Name/Alias to a UCS Image Table" on page 162.

If an alias name is supplied, it is used to schedule a printer for SYSOUT data sets. If no alias is supplied, an installation-defined SYSOUT class or a printer routing code (specified with the DEST parameter of JCL) should be used to assign the data set to the correct printer.

ADDING A UCS IMAGE TO THE IMAGE LIBRARY

Using the assembler and linkage editor, you may add a UCS image to those that reside in SYS1.IMAGELIB. No executable code is generated; the assembler prepares DCs, and the linkage editor puts them into SYS1.IMAGELIB. The new UCS image must be structured according to the following rules:

1. The member name must be 5 to 8 characters long; the first 4 characters must be the appropriate UCS prefix, as shown below.

UCS1 - 1403 Printer

UCS2 - 3211 Printer (or 3211 compatible printer)

UCS3 - 3203 Printer

These first four characters must be followed by a character set code, one to four characters long. Any valid combination of letters and numbers under assembler language rules is acceptable. However, the single letters U or C must not be used, because they are symbols for special conditions recognized by the system. The assigned character set code must be specified on the DD statement or SETPRT macro to load the image into the UCS buffer.

You can supply an alias name for a new image with the ALIAS statement of the linkage editor. (For more information on the ALIAS statement, see Linkage Editor and Loader.)

2. The first byte of the character set image load module specifies whether the image is a default. (Default images may be used by the system for jobs that do not request a specific image.) Specify the following in the first byte:

For JES2:

X'80' indicates a default image.

X'40' indicates that the output is to be folded.

X'C0' indicates default image and folding.

X'00' indicates that the image is not to be used as a default.

For non-JES2:

X'80' indicates a default image.

X'00' indicates that the image is not to be used as a default.

3. The second byte of the load module indicates the number of lines (n) to be printed for image verification. See "Verifying the UCS Image" on page 162 for more information on image verification.
4. Each byte of the next n bytes indicates the number of characters to be printed on each verification line. For the 3211 Printer, the maximum number of characters printed per line is 48; the bytes of associative bits (see note 5) are not printed during verification.
5. The UCS image itself must follow the previously described fields. The image must fill the number of bytes required by the printer; see the table below for image lengths. Note that, because of assembler language syntax, two apostrophes or two ampersands must be coded to represent a single apostrophe or a single ampersand, respectively, within a character set image.

Printer	Image Length
1403	240 bytes
3203	304 bytes (240 characters followed by 64 bytes of associative bits)
3211	512 bytes (432 characters followed by 15 bytes of X'00', 64 bytes of associative bits, and one reserved byte of X'00')

Associative bits must be coded to prevent data checks when adding a UCS image to SYS1.IMAGELIB. See the appropriate printer manual for more information on coding associative bits.

UCS CODING EXAMPLES

- Figure 29 contains an example of adding a 1403 UCS image, YN, to SYS1.IMAGELIB or the image library. Notes follow Figure 31 on page 161.
- Figure 30 on page 160 shows the code used to add a 3203 UCS image, YN, to SYS1.IMAGELIB or the image library.
- Figure 31 on page 161 shows the code used to add a 3211 UCS image, All, to SYS1.IMAGELIB or the image library.

```

//ADDYN      JOB  MSGLEVEL=1
//STEP      EXEC PROC=ASMFCL,PARM.ASM='NODECK,LOAD',
//          PARM.LKED='LIST,OL,REFR,RENT,XREF'
//ASM.SYSIN DD  *
UCSIYN      CSECT
            DC  X'80'          (THIS IS A DEFAULT IMAGE)
            DC  ALL(6)         (NUMBER OF LINES TO BE PRINTED)
            DC  ALL(39)        (39 CHARACTERS TO BE PRINTED ON LINE 1)
            DC  ALL(42)        (42 CHARACTERS TO BE PRINTED ON LINE 2)
            DC  ALL(39)        (39 CHARACTERS TO BE PRINTED ON LINE 3)
            DC  ALL(39)        (39 CHARACTERS TO BE PRINTED ON LINE 4)
            DC  ALL(42)        (42 CHARACTERS TO BE PRINTED ON LINE 5)
            DC  ALL(39)        (39 CHARACTERS TO BE PRINTED ON LINE 6)
*          THE FOLLOWING SIX LINES REPRESENT THE TRAIN IMAGE
            DC  C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.'
            DC  C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.#-$'
            DC  C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.'
            DC  C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.#-$'
            DC  C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.'
            DC  C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.#-$'
            END
/*
//LKED.SYSLMOD DD  DSNAMESYS1.IMAGELIB(UCSIYN),DISP=OLD,
//          SPACE=          (OVERRIDE SECONDARY ALLOCATION)

```

Figure 29. Sample Code to Add a 1403 UCS Image to SYS1.IMAGELIB

```

//ADYN3203      JOB  MSGLEVEL=1
//STEP          EXEC PROC=ASMFCL, PARM.ASM='NODECK, LOAD',
//              PARM.LKED='LIST, OL, REFR, RENT, XREF'
//ASM.SYSIN     DD   *
UCS3YN          CSECT
                DC   X'80'          (THIS IS A DEFAULT IMAGE)
                DC   AL1(6)         (NUMBER OF LINES TO BE PRINTED)
                DC   AL1(39)        (39 CHARACTERS TO BE PRINTED ON LINE 1)
                DC   AL1(42)        (42 CHARACTERS TO BE PRINTED ON LINE 2)
                DC   AL1(39)        (39 CHARACTERS TO BE PRINTED ON LINE 3)
                DC   AL1(39)        (39 CHARACTERS TO BE PRINTED ON LINE 4)
                DC   AL1(42)        (42 CHARACTERS TO BE PRINTED ON LINE 5)
                DC   AL1(39)        (39 CHARACTERS TO BE PRINTED ON LINE 6)
*              THE FOLLOWING SIX LINES REPRESENT THE TRAIN IMAGE
                DC   C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.'
                DC   C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.#-$'
                DC   C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.'
                DC   C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.#-$'
                DC   C'1234567890STABCDEFHIJKLMNOPQRSTUVWXYZ*,.'
*              THE FOLLOWING FOUR DC INSTRUCTIONS DEFINE THE ASSOCIATIVE BITS,
*              UCSB BYTE POSITIONS 241-304
                DC   X'C01010101010101010101000400000000000010'
                DC   X'101010101010101010101000404000000040001010'
                DC   X'10101010101010004000000000101010101010'
                DC   X'10101010004000000000'
                END
/*
//LKED.SYSLMOD DD  DSNAME=SYS1.IMAGELIB(UCS3YN), DISP=OLD,
//                SPACE=          (OVERRIDE SECONDARY ALLOCATION)

```

Figure 30. Sample Code to Add a 3203 UCS Image to SYS1.IMAGELIB

```

//ADDA11      JOB  MSGLEVEL=1
//STEP        EXEC  PROC=ASMFCL, PARM.ASM='NODECK, LOAD',
//            PARM.LKED='LIST, OL, REFR, RENT, XREF'
//ASM.SYSIN   DD   *
UCS2A11      CSECT
DC  X'80'      (THIS IS A DEFAULT IMAGE)
DC  AL1(9)    (NUMBER OF LINES TO BE PRINTED)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 1)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 2)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 3)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 4)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 5)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 6)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 7)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 8)
DC  AL1(48)   (48 CHARACTERS TO BE PRINTED ON LINE 9)
*           THE FOLLOWING NINE LINES REPRESENT THE TRAIN IMAGE
*           NOTE 2 AMPERSANDS MUST BE CODED TO GET 1 IN ASSEMBLER SYNTAX
DC  C'1<.+IHGFEDCBA*$$-RQPONMLKJ%,&&ZYXWVUTS/@#098765432'
DC  15X'00'   (RESERVED FIELD, BYTES 433-447)
*           THE FOLLOWING FOUR DC INSTRUCTIONS DEFINE THE ASSOCIATIVE BITS,
*           UCSB BYTE POSITIONS 448-511
DC  X'C0101010101010101010100040404240004010'
DC  X'1010101010101010101000404041000040401010'
DC  X'10101010101010004040000000101010101010'
DC  X'10101010004040444800'
DC  X'00'     (RESERVED FIELD, BYTE 512)
END
/*
//LKED.SYSLMOD DD  DSNAME=SYS1.IMAGELIB(UCS2A11),DISP=OLD,
//                SPACE=          (OVERRIDE SECONDARY ALLOCATION)

```

Figure 31. Sample Code to Add a 3211 UCS Image to SYS1.IMAGELIB

Notes to Figure 29 on page 159, Figure 30 on page 160, and Figure 31:

1. The RENT and REFR linkage editor attributes are used for performance considerations in a paging environment. They are required attributes.
2. For the 3203 and 3211 Printers, the 64 bytes of associative bits must be coded to avoid data checks. To determine how to code these bits for a particular image, see IBM 3203 Printer Component Description and Operator's Guide or IBM 3211 Printer, 3216 Interchangeable Train Cartridge, and 3811 Printer Control Unit Component Description and Operator's Guide.
3. Executing the ASMFCL procedure does not actually generate executable code. The assembler/linkage editor is used to place the UCS image into SYS1.IMAGELIB.
4. The SPACE parameter is overridden here because the ASMFCL cataloged procedure has secondary allocation specified. Elimination of the override causes the original secondary allocation amount to be used.

| Verifying the UCS Image

For the 1403 (with the UCS feature), 3203, 3211, 3262, and 4245 Printers, the UCS image can be displayed on the printer for visual verification using either of the following parameters:

- In JCL: UCS=(character set code,VERIFY)
- In the SETPRT macro: UCS=(character set code,V)

These parameters can also be used for the 4248 Printer.

Because the UCS image cannot be read directly from the 3262 Model 5 and 4248, only the header information on the verification display will be printed out. The verification display header appears in the format shown below.

UCS IMAGE VERIFICATION image id [,FOLD] [description]

image id

The 1- to 4-character name of the UCS image.

description

The descriptive information supplied for this UCS image in the UCS image table.

The 4245 also, optionally, prints the image.

For more information about the UCS VERIFY parameters, see JCL and Data Administration: Macro Instruction Reference.

| ADDING A UCS IMAGE NAME/ALIAS TO A UCS IMAGE TABLE

SYS1.IMAGELIB does not contain UCS images for the 3262 Model 5, 4245, or 4248 Printers. Instead, the image for each band is stored in the printer, and automatically loaded into the UCS buffer when the machine is powered on or a new band is installed. Information about these images is recorded in the IBM-supplied UCS image table, which resides in SYS1.IMAGELIB.

| UCS Image Table Structure

SYS1.IMAGELIB contains one UCS image table for each type of printer that supports image tables. The image table contains an entry for most installation-standard IBM-supplied bands. For the 4245, the table is called UCS5. For the 4248 and the 3262 Model 5, the table is called UCS6. A typical UCS image table entry is shown in Figure 32 on page 163.

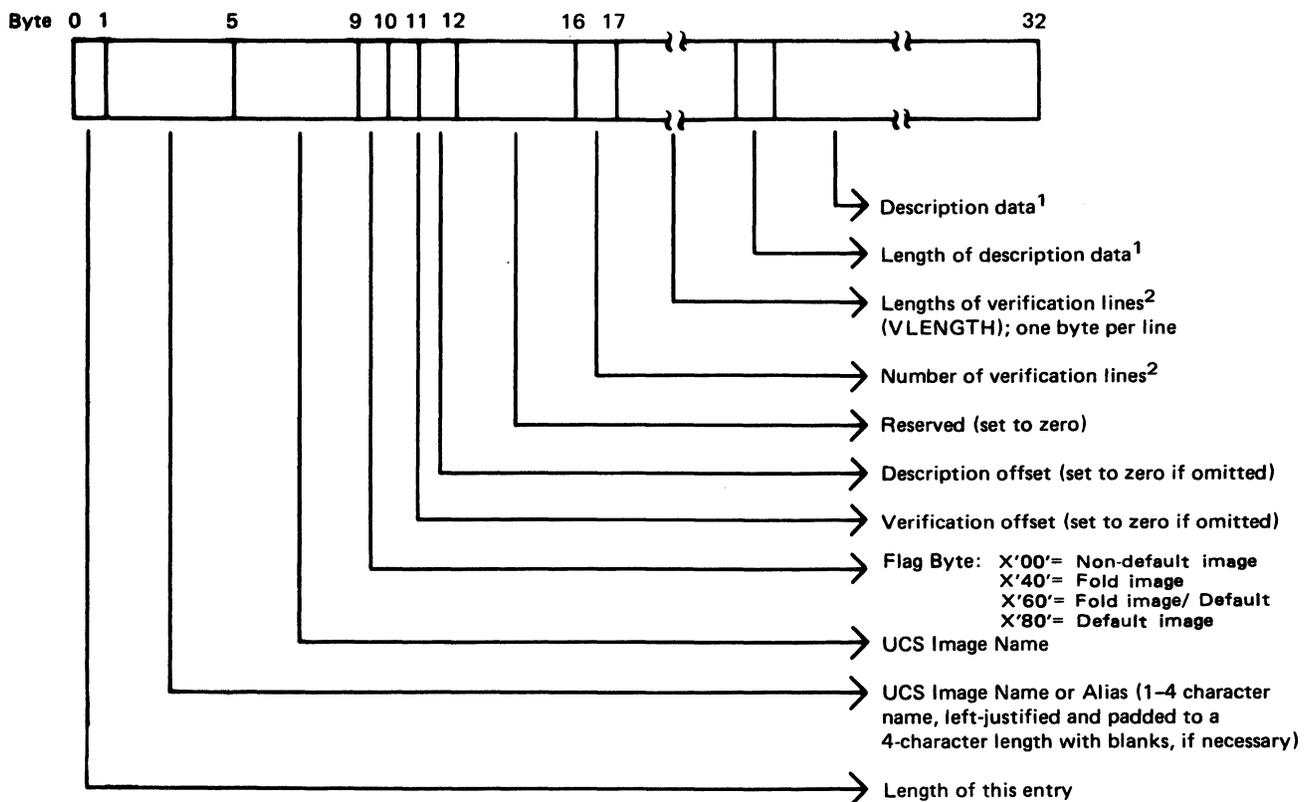


Figure 32. UCS Image Table Entry Format

Notes to Figure 32:

1. This field is optional.
2. This field is optional for the 4245 Printer. For the 4248, this field does not apply and is set to X'00'.

The contents of the UCS image table UCS5 (IGGUCS5 macro), for the 4245 Printer, are shown in Figure 33.

Name	Alias	Default	Description
AN21	AN21	YES	Default UCS image
AN21	AN	NO	1403/3203 AN image
AN21	A11	NO	3211 A11 image
AN21	40E1	NO	4248 40E1 image
HN21	HN21	NO	Nondefault UCS image
HN21	HN	NO	1403/3203 HN image

Figure 33 (Part 1 of 2). UCS5 Image Table Contents

Name	Alias	Default	Description
HN21	H11	NO	3211 H11 image
HN21	4101	NO	4248 4101 image
PL21	PL21	NO	Nondefault UCS image
PL21	PN	NO	1403/3203 PN image
PL21	P11	NO	3211 P11 image
PL21	4121	NO	4248 4121 image
SN21	SN21	NO	Nondefault UCS image
SN21	4201	NO	4248 4201 image
TN21	TN21	NO	Nondefault UCS image
TN21	TN	NO	1403/3203 TN image
TN21	T11	NO	3211 T11 image
TN21	4181	NO	4248 4181 image
GN21	GN21	NO	Nondefault UCS image
GN21	G11	NO	3211 G11 image
GN21	41C1	NO	4248 41C1 image
RN21	RN21	NO	Nondefault UCS image
RN21	RN	NO	1403/3203 RN image
KA21	KA21	NO	Nondefault UCS image
KA21	4041	NO	4248 4041 image
KA22	KA22	NO	Nondefault UCS image
FC21	FC21	NO	Nondefault UCS image
FC21	4161	NO	4248 4161 image

Figure 33 (Part 2 of 2). UCS5 Image Table Contents

The contents of the UCS image table UCS6 (IGGUCS6 macro), for the 4248 Printer, are shown in Figure 34.

Name	Alias	Default	Description
40E1	40E1	YES	Default UCS image
40E1	AN21	NO	4245 AN21 image
40E1	AN	NO	1403/3203 AN image
40E1	A11	NO	3211 A11 image
4101	4101	NO	Nondefault UCS image
4101	HN21	NO	4245 HN21 image
4101	HN	NO	1403/3203 HN image
4101	H11	NO	3211 H11 image
41C1	41C1	NO	Nondefault UCS image
41C1	GN21	NO	4245 GN21 image
41C1	G11	NO	3211 G11 image
4121	4121	NO	Nondefault UCS image
4121	PL21	NO	4245 PL21 image
4121	PN	NO	1403/3203 PN image
4121	P11	NO	3211 P11 image
4181	4181	NO	Nondefault UCS image
4181	TN21	NO	4245 TN21 image
4181	TN	NO	1403/3203 TN image
4181	T11	NO	3211 T11 image
4061	4061	NO	Nondefault UCS image
40C1	40C1	NO	Nondefault UCS image
4161	4161	NO	Nondefault UCS image
4161	FC21	NO	4245 FC21 image
4201	4201	NO	Nondefault UCS image
4201	SN21	NO	4245 SN21 image
4041	4041	NO	Nondefault UCS image
4041	KA21	NO	4245 KA21 image

Figure 34. UCS6 Image Table Contents

Note: The image tables for the 4245 and 4248 Printers include USA and Canada band IDs only. To support other national band IDs, you must modify the UCS image table. See "Adding/Modifying a UCS Image Table Entry" on page 166.

The 3262 Model 5 Printer uses the 4248 UCS image table, UCS6. However, no 3262 Model 5 band names or aliases are provided by IBM in UCS6. In order to use 3262 Model 5 UCS images, you must add the names and aliases to UCS6 yourself. "Adding/Modifying a UCS Image Table Entry" on page 166 describes how to add entries

to the UCS image table. For a list of the bands available for the 3262 Model 5, see IBM 3262 Printer Model 5 Product Description.

Adding/Modifying a UCS Image Table Entry

If you plan to use a new UCS image name/alias with the 3262 Model 5, 4245, or 4248 Printer, you must add an entry for that image name/alias to the appropriate UCS image table. Similarly, if you want to select a new default image or change the description on an old image, you must make the change in the image table.

To build new UCS table entries, or to change the format of old entries, use the following procedure:

1. Issue the IGGUCSIT macro, as described below, to build a new UCS image table entry. A new entry is built even if it is intended to replace an existing entry supplied by IBM. Because the new entry is found first, the previous entry is never found and thus is effectively replaced.
2. Include the UCS image table, source using the IGGUCS5 or IGGUCS6 macro, both of which are found in SYS1.MACLIB.
3. Reassemble the image table module (UCS5 or UCS6).
4. Link-edit the reassembled module into SYS1.IMAGELIB.

The IGGUCSIT macro instruction has the following format:

IGGUCSIT	MF={LIST DSECT} ,NAME= <u>image name</u> [,ALIAS=image alias] [,DEFAULT={YES NO}] [,DESCR=description] [,DEVICE={4245 4248}] [,VLENGTH=(n1,n2,. . .n)]
-----------------	---

MF={LIST|DSECT}

specifies the form of the macro instruction.

LIST

produces a UCS image table entry based on the information supplied in other IGGUCSIT parameters. If LIST is selected or allowed to default, the NAME parameter must also be coded.

DSECT

produces a DSECT for a single UCS image table entry, similar to the sample entry shown in Figure 32 on page 163. If you code DSECT, all other parameters of IGGUCSIT are ignored.

LIST is the default.

NAME=image name

specifies the 1 to 4 character UCS image name.

ALIAS=image alias

specifies a 1 to 4 character alias name for the UCS image. If ALIAS is not specified, the image name coded in the NAME parameter will be entered in the UCS image table.

DEFAULT={YES|NO}

indicates whether the new UCS image is to be used as a default value.

YES

indicates that this UCS image is a default. Default images are used by the system for jobs that do not request a specific image.

NO indicates that this UCS image should not be used as a default.

If the DEFAULT parameter is not specified, the new UCS image is not used as a default.

DESCR=description
specifies descriptive information about the new UCS image. description can be up to 32 EBCDIC or hexadecimal characters long. You cannot use EBCDIC and hexadecimal characters in combination.

Descriptive information is placed in the header line of the verification display, following the real UCS image name. If you omit the DESCR parameter, no description appears in the display. For more information on the verification display, see "Verifying the UCS Image" on page 162.

If VLENGTH is not specified for the 4245 Printer, the DESCR parameter is ignored.

DEVICE={4245|4248}
specifies the type of device for which an image table entry is to be created.

If you specify MF=LIST on the first invocation of the IGGUCSIT macro, DEVICE defaults to 4245. The default for subsequent invocations is the printer type that you specified (or the default) on the first invocation. Table entries with different DEVICE specifications are not allowed.

For the 3262 Model 5 Printer, DEVICE=4248 should be specified in order to create the appropriate form of the image table entry.

VLENGTH=(n1,n2,...n)
specifies the length(s) of each line in the UCS verification display. The length of each line must be specified separately, even if all lines are of the same length.

n1 is the length of print line 1; n2 is the length of print line 2; n is the length of the last print line. To display the complete image, the sum of the verification line lengths should equal 350.

For details on the verification report, see "Verifying the UCS Image" on page 162.

The VLENGTH parameter is not valid for the 3262 Model 5 or 4248 Printer.

EXAMPLES OF ADDING TO THE UCS IMAGE TABLE

Example 1: Adding a New Band ID to the 4245 UCS Image Table (UCS5)

In this example, the band name RPQ1 with description "RPQ BAND" is added to UCS5. In the UCS verification display, 7 lines of 50 characters each are printed. Macro IGGUCS5 causes the UCS image table source (as distributed by IBM) to be included in the table entry.

```

//UCS5      JOB      . . . . .
//          EXEC    ASMFCL,
//          PARM.ASM='NODECK,LOAD',
//          PARM.LKED='OL,RENT,REUS'
//SYSPRINT DD      SYSOUT=A
//ASM.SYSIN DD     *
//          TITLE  'UPDATED UCS5 IMAGE TABLE'
UCS5  CSECT
      IGGUCSIT  NAME=RPQ1,
                VLENGTH=(50,50,50,50,50,50,50),
                DESCR='RPQ BAND'
      IGGUCS5
      END
/*
//LKED.SYSLMOD DD DSN=SYS1.IMAGELIB(UCS5),DISP=OLD,
//                SPACE=      (OVERRIDE SECONDARY ALLOCATION)

```

72

X
X

Notes to Example 1:

1. The RENT and REUS linkage editor attributes are used for performance considerations in a paging environment. They are required attributes.
2. Executing the ASMFCL procedure does not actually generate executable code. The assembler/linkage editor is used to place the UCS image table entry into SYS1.IMAGELIB.
3. The SPACE parameter is overridden here because the ASMFCL cataloged procedure has secondary allocation specified. Elimination of the override causes the original secondary allocation amount to be used.

Example 2: Adding a New Default Entry to the 4248 UCS Image Table (UCS6).

In the following example, the band name 40E1 with description "40E1 DEFAULT BAND" is added to UCS6 and defined as a default band. An alias name, HN21, is also defined for band 40E1. Macro IGGUCS6 causes the UCS image table source (as distributed by IBM) to be included in the table entry.

```

//UCS6      JOB      . . . . .
//          EXEC    ASMFCL,
//          PARM.ASM='NODECK,LOAD',
//          PARM.LKED='OL,RENT,REUS'
//SYSPRINT DD      SYSOUT=A
//ASM.SYSIN DD     *
//          TITLE  'UPDATED UCS6 IMAGE TABLE'
UCS6  CSECT
      IGGUCSIT  NAME=40E1,
                DEVICE=4248,
                ALIAS=HN21,
                DEFAULT=YES,
                DESCR='40E1 DEFAULT BAND'
      IGGUCS6
      END
/*
//LKED.SYSLMOD DD DSN=SYS1.IMAGELIB(UCS6),DISP=OLD,
//                SPACE=      (OVERRIDE SECONDARY ALLOCATION)

```

72

X
X
X
X

Note that this method creates a duplicate entry for 40E1 that becomes the first entry in the table. Because the table is

searched sequentially, the new entry is always found before the old entry, thus effectively replacing the old entry.

Notes to Example 2:

1. The RENT and REUS linkage editor attributes are used for performance considerations in a paging environment. They are required attributes.
2. Executing the ASMFCL procedure does not generate executable code. The assembler/linkage editor places the UCS image table entry into SYS1.IMAGELIB.
3. The SPACE parameter is overridden because the ASMFCL cataloged procedure has secondary allocation specified. Elimination of the override causes the original secondary allocation amount to be used.

ADDING AN FCB IMAGE TO THE IMAGE LIBRARY

Two standard FCB images, STD1 and STD2, are included in SYS1.IMAGELIB during system generation for the following printers:

3203
3211
4245
3262 Model 5
4248

The 4248 and 3262 Model 5 Printers also accept FCBs that can be used with the 3203, 3211, and 4245 Printers. (These are referred to as 3211 format FCBs.)

STD1 sets line spacing at 6 lines per inch for an 8-1/2 inch form; STD2 is a default FCB image that sets line spacing at 6 lines per inch for an 11-inch form. Channels for both images are evenly spaced, with Channel 1 on the fourth line and Channel 9 on the last line. See Figure 35 on page 170 and Figure 36 on page 170 for sample STD1 and STD2 images.

The 3262 Model 5, the 4245, or the 4248 Printer loads a default FCB image into the buffer when the machine is powered on. The 3262 Model 5 default FCB image is an 11-inch form with 6 lines per inch, a Channel 1 on the third print line, and a Channel 12 on line 64. The 4245 default FCB image is an 11-inch form with 6 lines per inch and a Channel 1 on the first print line. The 4248 default FCB image is the last FCB image loaded.

You should use the IEBIMAGE utility to create and modify FCB modules for the 3800 Printing Subsystem. You should also use it to create and modify FCB images for the 3262 Model 5 or 4248 Printer (4248 format FCBs). For information on IEBIMAGE and the format of the 4248 FCB image, see Utilities.

```

FCB2STD1  CSECT
DC        X'80'          DEFAULT
DC        AL1(48)       FCB IMAGE LENGTH = 48
DC        X'000000'     LINE 1, 2, 3
DC        X'01'         LINE 4, CHANNEL 1
DC        X'000000'     LINE 5, 6, 7
DC        X'02'         LINE 8, CHANNEL 2
DC        X'000000'     LINE 9, 10, 11
DC        X'03'         LINE 12, CHANNEL 3
DC        X'000000'     LINE 13, 14, 15
DC        X'04'         LINE 16, CHANNEL 4
DC        X'000000'     LINE 17, 18, 19
DC        X'05'         LINE 20, CHANNEL 5
DC        X'000000'     LINE 21, 22, 23
DC        X'06'         LINE 24, CHANNEL 6
DC        X'000000'     LINE 25, 26, 27
DC        X'07'         LINE 28, CHANNEL 7
DC        X'000000'     LINE 29, 30, 31
DC        X'08'         LINE 32, CHANNEL 8
DC        X'000000'     LINE 33, 34, 35
DC        X'0A'         LINE 36, CHANNEL 10
DC        X'000000'     LINE 37, 38, 39
DC        X'0B'         LINE 40, CHANNEL 11
DC        X'000000'     LINE 41, 42, 43
DC        X'0C'         LINE 44, CHANNEL 12
DC        X'000000'     LINE 45, 46, 47
DC        X'19'         LINE 48, CHANNEL 9-END OF FCB IMAGE
END

```

Figure 35. Sample of the Standard FCB Image STD1

```

FCB2STD2  CSECT
DC        X'80'          DEFAULT
DC        AL1(66)       FCB IMAGE LENGTH = 66
DC        X'000000'     LINE 1, 2, 3
DC        X'01'         LINE 4, CHANNEL 1
DC        X'0000000000'  LINE 5, 6, 7, 8, 9
DC        X'02'         LINE 10, CHANNEL 2
DC        X'0000000000'  LINE 11, 12, 13, 14, 15
DC        X'03'         LINE 16, CHANNEL 3
DC        X'0000000000'  LINE 17, 18, 19, 20, 21
DC        X'04'         LINE 22, CHANNEL 4
DC        X'0000000000'  LINE 23, 24, 25, 26, 27
DC        X'05'         LINE 28, CHANNEL 5
DC        X'0000000000'  LINE 29, 30, 31, 32, 33
DC        X'06'         LINE 34, CHANNEL 6
DC        X'0000000000'  LINE 35, 36, 37, 38, 39
DC        X'07'         LINE 40, CHANNEL 7
DC        X'0000000000'  LINE 41, 42, 43, 44, 45
DC        X'08'         LINE 46, CHANNEL 8
DC        X'0000000000'  LINE 47, 48, 49, 50, 51
DC        X'0A'         LINE 52, CHANNEL 10
DC        X'0000000000'  LINE 53, 54, 55, 56, 57
DC        X'0B'         LINE 58, CHANNEL 11
DC        X'0000000000'  LINE 59, 60, 61, 62, 63
DC        X'0C'         LINE 64, CHANNEL 12
DC        X'00'         LINE 65
DC        X'19'         LINE 66, CHANNEL 9-END OF FCB IMAGE
END

```

Figure 36. Sample of the Standard FCB Image STD2

You may add a 3211 format FCB image to those that reside in SYS1.IMAGELIB, using the assembler and linkage editor. No executable code is generated; the assembler prepares DCs, and the linkage editor puts them into SYS1.IMAGELIB. The new FCB image must be structured according to the following rules:

1. The member name may not exceed eight bytes. The first four characters of the name must be FCB2. The characters that follow identify the FCB image and are referred to as the "image identifier" (ID). Any combination of valid assembler language characters can be used, with the exception of a single "C" or "U", because these are used by the system to recognize special conditions. The image identifier must be specified in the FCB keyword of a DD statement or in the SETPRT macro to load the image into the FCB buffer.
2. The first byte of the FCB load module specifies whether the image is a default. (Default images may be used by the system for jobs that do not request a specific image.) Specify the following in the first byte:
 - X'80' indicates a default image
 - X'00' indicates a nondefault image
3. The second byte of the load module indicates the number of bytes to be transferred to the control unit to load the FCB image. This count includes the byte, if used, for the print position indexing feature.
4. The third byte of the load module (the first byte of the FCB image) is either the print position indexing byte, or the lines-per-inch byte. The print position indexing byte is optional and, when used, precedes the lines-per-inch byte. The 4245 and 4248 Printers accept and discard the index byte if it is present, because neither printer supports the indexing feature. A description of the print position indexing feature and its use will be found in IBM 3211 Printer, 3216 Interchangeable Train Cartridge, and 3811 Printer Control Unit Component Description and Operator's Guide.

The special index flag in the third byte contains X'80' plus a binary index value, from 1 to 32 (the default is 1). This index value sets the left margin: 1 indicates flush-left; any other value indicates a line indented that many spaces.

The form image begins with the lines-per-inch (LPI) byte. The LPI byte defines the number of lines per inch (6 or 8) and also represents the first line of the page. It may or may not also contain a channel identifier.

Typically, the length of an FCB image is consistent with the length of the form it represents. For example, an 8-1/2 inch form to be printed at 6 LPI has an FCB image that is 51 bytes long (8-1/2 inches times 6 LPI).

The LPI byte appears as follows:

- X'1n' sets 8 LPI
- X'0n' sets 6 LPI

5. All remaining bytes (lines) must contain X'0n', except the last byte, which must be X'1n'. The letter n can be a hexadecimal value from 1 to C, representing a channel (one to 12), or it can be 0, which means no channel is indicated.

In Figure 37 on page 172, an FCB load module is assembled and added to SYS1.IMAGELIB. The image defines a print density of 8 lines per inch on an 11-inch form, with a right shift of 15 line character positions (1-1/2 inches).

```

//ADDFCB      JOB MSGLEVEL=1
//STEP        EXEC PROC=ASMFCL, PARM.ASM='NODECK,LOAD',
//            PARM.LKED='LIST,OL,REFR,RENT,XREF'
//ASM.SYSIN   DD *
FCB2ID1      CSECT
*THIS EXAMPLE IS FOR A FORM LENGTH OF 11 INCHES WITH 8 LPI (88 LINES)
DC X'80'      THIS IS A DEFAULT IMAGE
DC AL1(89)    LENGTH OF FCB IMAGE AND INDEXING BYTE
DC X'8F'      OFFSET 15 CHARACTERS TO THE RIGHT
DC X'10'      8 LINES PER INCH-NO CHANNEL FOR LINE 1
DC XL4'0'     4 LINES NO CHANNEL
DC X'01'      CHANNEL 1 IN LINE 6
DC XL6'0'     6 LINES NO CHANNEL
DC X'02'      CHANNEL 2 IN LINE 13
DC XL6'0'     6 LINES NO CHANNEL
DC X'03'      CHANNEL 3 IN LINE 20
DC XL6'0'     6 LINES NO CHANNEL
DC X'04'      CHANNEL 4 IN LINE 27
DC XL6'0'     6 LINES NO CHANNEL
DC X'05'      CHANNEL 5 IN LINE 34
DC XL6'0'     6 LINES NO CHANNEL
DC X'06'      CHANNEL 6 IN LINE 41
DC XL6'0'     6 LINES NO CHANNEL
DC X'07'      CHANNEL 7 IN LINE 48
DC XL6'0'     6 LINES NO CHANNEL
DC X'08'      CHANNEL 8 IN LINE 55
DC XL6'0'     6 LINES NO CHANNEL
DC X'09'      CHANNEL 9 IN LINE 62
DC XL6'0'     6 LINES NO CHANNEL
DC X'0A'      CHANNEL 10 IN LINE 69
DC XL6'0'     6 LINES NO CHANNEL
DC X'0B'      CHANNEL 11 IN LINE 76
DC XL6'0'     6 LINES NO CHANNEL
DC X'0C'      CHANNEL 12 IN LINE 83
DC XL4'0'     4 LINES NO CHANNEL
DC X'10'      POSITION 88 LAST LINE IN IMAGE
END
/*
//LKED.SYSLMOD DD DSNAME=SYS1.IMAGELIB(FCB2ID1),DISP=OLD,
//              SPACE= (OVERRIDE SECONDARY ALLOCATION)

```

Figure 37. Sample Code to Assemble and Add an FCB Load Module to SYS1.IMAGELIB

Notes to Figure 37:

1. The RENT and REFR linkage editor attributes are used for performance considerations in a paging environment. They are required attributes.
2. Executing the ASMFCL procedure does not actually generate executable code. The assembler/linkage editor is used to place the FCB image into SYS1.IMAGELIB.
3. The SPACE parameter is overridden here because the ASMFCL cataloged procedure has secondary allocation specified. Elimination of the override causes the original secondary allocation amount to be used.

RETRIEVING AN FCB IMAGE FROM SYS1.IMAGELIB

If you want to modify an FCB image in virtual storage before loading it into a forms control buffer, you can use this sequence of macro instructions to read the FCB image into virtual storage.

1. An IMGLIB macro instruction, along with the OPEN parameter
2. A BLDL macro instruction to determine whether the FCB image you want is in the image library
3. A LOAD macro instruction to load the image into virtual storage

After the image has been read in, you should issue the IMGLIB macro instruction with the CLOSE parameter and the address of the DCB that was built by the first IMGLIB macro. A SETPRT macro instruction can be used to load the forms control buffer with the modified image. Printers other than the 3800 will require the use of an FCB entry in an exit list, as described in Data Administration Guide.

The format of the BLDL and SETPRT macros is given in Data Administration: Macro Instruction Reference; the format of the LOAD macro is given in Supervisor Services and Macro Instructions.

The format of the IMGLIB macro is shown below:

[symbol]	IMGLIB	{OPEN CLOSE,addr}
----------	--------	-------------------

OPEN

specifies that a DCB is to be built for SYS1.IMAGELIB and that SYS1.IMAGELIB is to be opened. The address of the DCB is returned in register 1.

CLOSE

specifies that SYS1.IMAGELIB is to be closed.

addr

specifies the RX-type address of the word that points to the DCB. If coded in the form (reg), the register in parentheses then contains the address of the DCB, not the address of the fullword.

Return codes from the IMGLIB OPEN macro are shown below:

Return Code	Meaning
0(X'00')	Operation successful.
4(X'04')	Either the volume containing SYS1.IMAGELIB is not mounted or a required catalog volume is not mounted.
8(X'08')	Either SYS1.IMAGELIB does not exist on the volume to which the catalog points, or SYS1.IMAGELIB is not cataloged.
12(X'0C')	An error occurred in reading the catalog or VTOC.

BLDL and LOAD are the only macros that may refer to the DCB built by the IMGLIB macro.

CHAPTER 8. JES2 SUPPORT FOR THE IBM 1403, 3203 MODEL 5, AND 3211 PRINTERS

UCS ALIAS NAMES

The system assigns an alias for each installation-standard print chain not actually defined on a given printer. This provides JES2 with flexibility in scheduling printers for SYSOUT data sets. For example, a request for the 1403 TN train would be assigned the T11 train if the data set were printed on a 3211. The assigned alias names that follow the naming conventions currently used in SYS1.IMAGELIB are:

Image	Alias
UCS1AN	UCS1A11
UCS1HN	UCS1H11
UCS1PN	UCS1P11
UCS1TN	UCS1T11
UCS2A11	UCS2AN
UCS2H11	UCS2HN
UCS2P11	UCS2PN,UCS2RN,UCS2QN
UCS2T11	UCS2TN

The image and alias names are included in SYS1.IMAGELIB at system generation.

Some trains, such as SN and G11, do not have aliases because neither has an equivalent train on the other printer. An installation can assign an alias, if it so chooses. (For details about the ALIAS statement, see Linkage Editor and Loader.) If an alias is supplied, JES2 will use it. If an alias is not supplied, an installation-defined SYSOUT class or a printer routing code (specified via the DEST parameter) should be used to assign the data set to the correct printer. If a SYSOUT class or a printer routing code is not used and if JES2 is directed to print a data set on a printer for which the proper image is not supplied, JES2 notifies the operator. The operator can then print the data set with a valid train or redirect the data set to the proper printer via the '\$E' command.

If an installation defines a new train, it can supply an alias name for that train, via the linkage editor ALIAS statement, when including the image in SYS1.IMAGELIB.

THE 3211 INDEXING FEATURE

JES2 supports the 3211 Indexing Feature in two ways:

1. Specification of the INDEX parameter on the /XOUTPUT card.
2. The extended FCB image:

JES2 supplies two special FCBs: FCB26 for 6 lines per inch and FCB28 for 8 lines per inch (specified as FCB=6 and FCB=8, respectively). These FCBs contain a channel 1 indication in position 1, a special index flag in the third byte, and the number of lines per inch in the fourth byte of the image.

The special index flag in the third byte of FCB26 and FCB28 contains X'80' plus a binary index value, in the range 1 to 32 (default=1). The index value sets the left margin (1 indicates flush-left position; other values cause indentation of the print line by N-1 positions).

If any other FCB images are to be used by JES2, they must specify channel 1 in position 1; otherwise, JES2 incorrectly positions the forms in the printer. (STD1 and STD2 do not specify channel 1 in position 1 and therefore must not be specified, unless altered, for JES2.)

If the third byte of any other FCB image contains a data character (specifying the number of lines per inch) other than X'80', JES2 uses that specification and supplies an index value of 1.

IBM 3203 MODEL 5 PRINTER

The IBM 3203 Model 5 Printer is treated as a 3211 Printer by JES2, except that the 3203 Model 5 does not support the 3211 indexing feature, and any indexing commands from JES2 are ignored by the 3203 Model 5. The 3203 Model 5 uses 3211 FCB images and its own unique UCS images. UCS images are listed in Installation System Generation.

CHAPTER 9. CATALOG, SCRATCH, AND RENAME DUMMY MODULES

The load modules for CATALOG (SVC 26), SCRATCH (SVC 29), and RENAME (SVC 30) contain as their entry points the dummy modules IGG026DU, IGG029DU, and IGG030DU, respectively. These dummy modules immediately pass control to the first processing module for their respective SVCs without performing any processing themselves. The CATALOG dummy module IGG026DU receives control from SVC 26 and immediately passes control to module IGC0002F. The SCRATCH dummy module IGG029DU receives control from SVC 29 and immediately passes control to module IGC0002I. The RENAME dummy module, IGG030DU, receives control from SVC 30 and immediately passes control to IGC00030.

The load module for SCRATCH(SVC29) also contains the dummy module IGG029DM. The SCRATCH dummy module IGG029DM receives control from IGG0290D when an error return code of 4 or 8 is indicated, and immediately passes control to the location pointed to by register 14.

If you require special processing either before or after SVC 26, 29, or 30, you replace the appropriate dummy module(s) with your own module(s). Your replacement modules must follow all the characteristics and programming conventions for SVC routines. For information on characteristics of SVC routines, programming conventions for SVC routines, writing SVC routines, and inserting SVC routines, see Supervisor Services and Macro Instructions. Your modules may replace IGG026DU, IGG029DU, IGG029DM, and IGG030DU in SYS1.AOSD0 prior to system generation, or you may replace the dummy modules in SYS1.LPALIB after system generation. Information on how to replace the dummy modules with your modules can be obtained from the appropriate link-edit step of the STAGE I system generation output. You may also obtain link-edit information from the STAGE I system generation macro SGIEC4DM in SYS1.AGENLIB. You may apply PTFs to CATALOG, SCRATCH, or RENAME with SMP without modifying your own versions of IGG026DU, IGG029DU, IGG029DM, and IGG030DU.

The prolog of each of the dummy modules contains register conventions and other information about these modules.

CHAPTER 10. SPECIFYING BUFFER NUMBERS FOR DASD DATA SETS

The BUFNO keyword in the DCB macro and the BUFNO subparameter of the DCB keyword in the DD statement determine how many buffers are allocated when accessing a partitioned or sequential data set using QSAM. The NCP Keyword in the DCB macro determines how many un-CHECKED READ or WRITE macro instructions are allowed when accessing a sequential or partitioned data set using BSAM; one buffer is used for each READ or WRITE macro instruction.

The sequential access method can construct a channel program to transfer as many as 30 buffers or 240000 bytes of data, whichever is less. If BUFNO or NCP is less than 30, no more than that number of buffers can be transferred with a single channel program.

BUFNO is defaulted in OPEN to five if it is not specified for a QSAM DCB; NCP is defaulted to one in OPEN if it is not specified. The QSAM access method manages buffers. The user program must manage buffers when it uses BSAM.

PERFORMANCE CONSIDERATIONS

Buffer number and block size influence the rate at which data can be transferred and the operating system overhead per block. The use of more buffers reduces (per block transferred) the EXCP and IOS overhead and the time waiting for the DASD device to seek to the requested cylinder and rotate to the requested record (device latency time). However, if more buffers are allocated than a program can effectively process, the virtual pages containing those buffers will be paged out, effectively adding to the system overhead for the job. A large number of buffers also cause a large amount of real storage to be allocated to the job while the data is being transferred.

A job in a low-performance group may get swapped out more frequently than a higher priority job. The number of buffers allocated for the job contributes to the number of pages that have to be swapped out.

Programs that access data sets with small block size (for example, 80) can easily make effective use of 30 buffers, which fit in, at most, two 4096-byte pages. The advantage of 30 buffers over the default of five buffers is great: one channel program versus six channel programs to transfer 30 blocks.

At the other end of the spectrum, usage of data sets with large blocking factors such as full-track blocking on 3350 or half-track blocking on 3380 can still be effective when only three or four buffers, rather than five or more, are specified. The slightly lower DASD performance and small increase in EXCP and IOS instruction costs should be more than offset by a reduction in paging or swapping in a constrained environment.

It can be seen that proper selection of buffer number can have a positive effect on the elapsed time of a job and the system overhead associated with the job. The DCB OPEN installation exit can use installation criteria for a default buffer number for QSAM DCBs (for a description of the OPEN installation exit, see "DCB OPEN Installation Exit" on page 100.) The NCP field of the DCB must be set by the program for BSAM DCBs.

APPENDIX A. CVAF - VTOC ACCESS MACROS

CVAFDIR MACRO

OVERVIEW OF THE CVAFDIR MACRO

For an indexed or nonindexed VTOC, the CVAFDIR macro may be used to:

- Read or write a DSCB by specifying the name of the data set it represents
- Read or write a DSCB by specifying its address

In addition, for an indexed VTOC, the macro may be used to:

- Read or write VTOC index records
- Read and retain in virtual storage the first high-level VIER, and VIERS used during an index search.
- Read and retain in virtual storage the space map VIRs
- Free VIRs retained in virtual storage

SYNTAX

[<u>label</u>]	CVAFDIR	ACCESS=READ WRITE RLSE [,DSN= <u>addr</u>] [,BUFLIST= <u>addr</u>] [,VERIFY=YES NO] [,UCB=(<u>reg</u>) DEB= <u>addr</u>] [,IOAREA=KEEP (KEEP, <u>addr</u>)NOKEEP (NOKEEP, <u>addr</u>)] [,MAPRCDS=YES (YES, <u>addr</u>) NO (NO, <u>addr</u>)] [,IXRCDS=KEEP (KEEP, <u>addr</u>) NOKEEP (NOKEEP, <u>addr</u>)] [,BRANCH=YES ¹ NO (YES,SUP) (YES,PGM)] [,MF= <u>I L (E,<u>addr</u>)]</u>
------------------	---------	---

¹The default is SUP if YES is coded.

ACCESS: READ OR WRITE A DSCB OR VIR(S), OR RELEASE BUFFER LISTS

When ACCESS is READ or WRITE, a single DSCB is accessed for an indexed or nonindexed VTOC, or one or more VIRs are accessed for an indexed VTOC.

ACCESS=READ

Specifies that a single DSCB or one or more VIR(s) are to be read into a buffer whose address is in a buffer list.

If the buffer list is for a DSCB, only one entry is used in the buffer list. The first entry with the skip bit set to zero and with a nonzero buffer address is used.

All VIR(s) whose buffer list entry has the skip bit off will be read into a buffer.

DSN and BUFLIST are required if ACCESS=READ for a DSCB buffer list.

ACCESS=WRITE

Specifies that a single DSCB or one or more VIRs are to be written from buffer(s) whose address is in a buffer list.

WRITE is permitted with BRANCH=NO only if the caller is authorized by APF.

DSN and BUFLIST are required if ACCESS=WRITE for a DSCB buffer list.

If any buffer list entry has its modified bit set, only those entries with the modified bit set will be written. If no modify bits are on, all VIRs will be written.

ACCESS=RLSE

Applies only to VIR buffer lists. It requests the release of one or more buffers in the VIR buffer list chain identified in the BUFLIST keyword, and the release of each buffer list for which all buffers are released.

DSN and BUFLIST are not required if ACCESS=RLSE.

Only buffers in the buffer list with the skip bit set to zero and with a nonzero buffer address are released. The buffer list is not released if any entry has the skip bit set to one.

For an indexed VTOC, if ACCESS=RLSE is coded, buffer lists and buffers pointed to by the BUFLIST keyword will be released, along with buffer lists supplied in the CVAF parameter list CVMRCDS and CVIRCDS fields. If the CVMRCDS or the CVIRCDS buffers are supplied in the BUFLIST field, either directly or indirectly through chaining, the keyword MAPRCDS=YES, IXRCDS=KEEP, or MAPRCDS=(NO,0), IXRCDS=(NOKEEP,0) must be coded to prevent CVAF from freeing the buffers more than once. If buffers are released, the CVAF parameter list field pointing to the buffer list will be updated.

DSN: SPECIFY THE NAME OF THE DSCB**DSN=addr**

DSN specifies the address of a 44-byte data set name of the DSCB to be accessed.

DSN is required if ACCESS=READ or WRITE and the request is to read or write a DSCB. If a 140-byte DSCB is specified:

- CVAF validity checks the storage location, but ignores the contents of the location.
- You must specify an argument that points to an extent within the VTOC.

BUFLIST: SPECIFY ONE OR MORE BUFFER LISTS**BUFLIST=addr**

The BUFLIST keyword contains the address of a buffer list used to read or write a DSCB or VIRs.

VERIFY: VERIFY THAT A DSCB IS A FORMAT-0 DSCB**VERIFY=YES**

CVAF will verify that the DSCB is a format-0 DSCB before writing the DSCB. The first four bytes of the key will be compared with binary zeros. If the key does not start with four bytes of zeros, the DSCB will not be written and an error code will be returned.

VERIFY=NO

CVAF will not test the key of the DSCB.

Note: VERIFY applies only when writing a 140-byte DSCB. VERIFY is ignored when a VIR is written.

UCB|DEB: SPECIFY THE VTOC TO BE ACCESSED**UCB=(reg)**

Supplies the address of the UCB for the unit whose VTOC is to be accessed. An unauthorized caller may not supply a UCB to CVAF.

Note: Code the address of the UCB parameter as register (2-12). Coding an RX-Type address here gives you unpredictable results.

DEB=addr

Specifies the address of a DEB opened to the VTOC you want to access. CVAF does not allow output requests to the VTOC or VTOC index if you specify the DEB subparameter. If you are not authorized, you cannot perform any asynchronous activity (such as EXCP, CLOSE, EOV) against the data set represented by the DEB because CVAF removes the DEB from the DEB table for the duration of the CVAF call. If you are not authorized (neither APF authorized nor in a system key), you must specify a DEB address, not a UCB, to CVAFDIR. See "Identifying the Volume" on page 26 for further details.

If you supply a previously obtained CVAF I/O area through the IOAREA keyword, you need not specify the UCB or DEB keyword. Otherwise you must specify either the UCB or DEB keyword. If you specify a UCB, the UCB address in the CVPL is overlaid by the UCB address in the I/O area.

If you supply both the UCB and DEB addresses in the CVPL, CVAF uses the DEB address and overlays the UCB address in the CVPL with the UCB address in the DEB.

IOAREA: KEEP OR FREE THE I/O WORK AREA**IOAREA=KEEP**

Specifies the CVAF I/O area associated with the CVAF parameter list is to be kept upon completion of the CVAF request. IOAREA=KEEP may be coded with BRANCH=NO only if the caller is authorized (APF or system key).

If IOAREA=KEEP is coded, the caller must issue CVAF with IOAREA=NOKEEP specified at some future time, whether or not any further VTOC access is required: for example, the recovery routine of the caller of CVAF.

Coding IOAREA=KEEP allows subsequent CVAF requests to be more efficient, as certain initialization functions can be bypassed. Neither DEB nor UCB need be specified when a previously obtained CVAF I/O area is supplied; neither can they be changed.

When IOAREA=KEEP is first issued, CVAF returns the CVAF I/O area in the CVAF parameter list (CVIOAR). Subsequent calls of CVAF may use that same parameter list, and CVAF will obtain its I/O area from the CVIOAR.

When processing on the current volume is finished, release all areas that were kept.

IOAREA=(KEEP,addr)

Provides the address of a previously obtained I/O area. If a different CVAF parameter list is used, the previously

obtained I/O area may be passed to CVAF by coding its address as the second parameter of the IOAREA keyword.

IOAREA=NOKEEP

Causes the work area to be freed upon completion of the CVAF request.

IOAREA=(NOKEEP,addr)

Causes a previously obtained work area to be freed upon completion of the CVAF request.

MAPRCDS: KEEP OR FREE MAPRCDS BUFFER LIST AND BUFFERS

This keyword applies to an indexed VTOC only and specifies the disposition of the MAPRCDS buffer list and buffers.

MAPRCDS=YES

Specifies that the buffer list and buffers are to be retained at the end of processing.

If no buffer list address is in the CVAF parameter list, CVAF will read the MAP VIRs into buffers it obtains. The buffer list that contains the address and RBAs of the VIRs can be accessed after processing from the CVAF parameter list field, CVMRCDS. The buffer list and VIR buffers are in your protect key: subpool 0 if you are not authorized; 229 if you are.

When processing on the current volume is finished, release all areas that were kept.

MAPRCDS=(YES,addr)

If YES is coded and the buffer list address (CVMRCDS in CVAF parameter list) is supplied, VIRs are not read.

The CVMRCDS buffer list used in CVAFDIR macro can be passed to another CVAF macro call through the MAPRCDS keyword.

If MAPRCDS=YES is coded for a nonindexed VTOC, the function is performed, but an error code will be returned.

MAPRCDS=NO

If MAPRCDS=NO is coded, all the buffers without the skip bit on in the buffer list whose address is in the CVMRCDS field of the CVPL will be freed. If all the buffers are freed, the buffer list will also be freed.

MAPRCDS=(NO,addr)

Causes buffer lists and buffers previously obtained by CVAF to be freed.

You must free buffer lists and buffers obtained by CVAF. This can be done in one of three ways:

- By coding MAPRCDS=NO on the CVAFDIR macro that obtained the buffers
- By coding MAPRCDS=NO on a subsequent CVAF macro
- By coding CVAFDIR ACCESS=RLSE and providing the address of the buffer list in the BUFLIST keyword

Note: You must enqueue the VTOC and reserve the unit to maintain the integrity of MAP records read.

IXRCDS: RETAIN VIERS IN VIRTUAL STORAGE

This keyword applies to indexed VTOCs only.

IXRCDS=KEEP

Specifies that VIERS read into storage are to be kept in virtual storage. The VIERS are retained even if processing cannot complete successfully. The CVAF parameter list in field CVIRCDS will have the address of a buffer list containing the VIR buffer addresses and RBAs of the VIERS read.

The index search function will dynamically update the buffer list and, when necessary, obtain additional buffer lists and chain them together.

If KEEP is specified and no buffer list is supplied to CVAF in the CVPL, CVAF will obtain a buffer list and buffers and read the first high-level VIER. The address of the buffer list is placed in the CVMICDS field of the CVPL. The first high-level VIER will be checked for the VXFHLV bit and to see whether the VXVISE bit is off.

The buffer list and VIR buffers are in your protect key. The subpool is 0 if you are not authorized; it is subpool 229 if you are.

If IXRCDS=KEEP is coded for a nonindexed VTOC, a request to read or write a DSCB will be performed, but an error code will be returned.

When processing on the current volume is finished, release all areas that were kept.

IXRCDS=(KEEP,addr)

The index records buffer list address from one CVAF request is being passed to this CVAF parameter list by specifying its address as the second parameter in the IXRCDS keyword.

IXRCDS=NOKEEP

If NOKEEP is coded, the VIERS that are accessed (if any) are not retained. Furthermore, the buffer list supplied in the CVIRCDS field in the CVAF parameter list is released, as are all buffers found in the buffer list. If the skip bit is set in any entry in the buffer list, the buffer and buffer list will not be freed.

IXRCDS=(NOKEEP,addr)

Specifies that previously accessed VIERS are not to be retained.

You must free buffer lists and buffers obtained by CVAF. This can be done in one of three ways:

- By coding IXRCDS=NOKEEP on the CVAFDIR macro that obtained the buffers
- By coding IXRCDS=NOKEEP on a subsequent CVAF macro
- By coding CVAFDIR ACCESS=RLSE and providing the address of the buffer list in the BUFLIST keyword

Note: You must enqueue the VTOC and reserve the unit to maintain the integrity of the VIERS read.

BRANCH: SPECIFY THE ENTRY TO THE MACRO

BRANCH=(YES,SUP)

Requests that the branch entry to CVAFDIR be used. You must be in supervisor state. Protect key checking is bypassed.

An 18-word save area must be supplied if BRANCH=YES is coded. No lock may be held on entry to CVAF. SRB mode is not allowed.

BRANCH=YES

Equivalent to BRANCH=(YES,SUP), because SUP is the default when YES is coded. Protect key checking is bypassed.

BRANCH=(YES,PGM)

Requests the branch entry. You must be authorized by APF and be in problem state. Protect key checking is bypassed.

BRANCH=NO

Requests the SVC entry. You must be authorized by APF if any output operations are requested. Protect key checking is performed.

MF: SPECIFY THE FORM OF THE MACRO

This keyword specifies whether the list, execute, or normal form of the macro is requested.

MF=I

If I is coded or if neither L nor E is coded, the CVAF parameter list is generated and CVAF is called. This is the normal form of the macro.

MF=L

L indicates the list form of the macro. A parameter list is generated, but CVAF is not called.

MF=(E,addr)

E indicates the execute form of the macro. The CVAF parameter list whose address is in 'addr' can be modified by this form of the macro.

RETURN CODES FROM THE CVAFDIR MACRO

On return from CVAF, register 1 contains the address of the CVPL (CVAF parameter list), and register 15 contains one of the following return codes:

Code	Meaning
00(X'00')	The request was successful. However, if the CVAFDIR request is to read or write a DSCB and a VTOC index structure error is encountered, the CVSTAT field indicates the structure error encountered. (CVSTAT code descriptions are in Appendix C, "VTOC Index Error Message and Associated Codes" on page 221.)
04(X'04')	An error occurred. The CVSTAT field in the CVPL contains an indication of the cause of the error. (CVSTAT code descriptions are in Appendix C, "VTOC Index Error Message and Associated Codes" on page 221.)
08(X'08')	Invalid VTOC index structure while processing a request to read or write a VTOC index record. The CVSTAT field in the CVPL contains an indication of the cause of the error. (CVSTAT code descriptions are in Appendix C, "VTOC Index Error Message and Associated Codes" on page 221.)

12(X'0C') The CVAF parameter list is not in your protect key or is invalid (the ID is invalid, or the length field is incorrect, or the CVFCTN field is invalid). The CVPL has not been modified.

16(X'10') An I/O error was encountered.

CVAFDSM MACRO

OVERVIEW OF THE CVAFDSM MACRO

The CVAFDSM macro may be used for an indexed VTOC to:

- Obtain one or more extents that describe unallocated space on the volume
- Obtain a count of free DSCBs on the VTOC
- Obtain a count of free VTOC index records in the VTOC index.

SYNTAX

[label]	CVAFDSM	ACCESS=MAPDATA ,MAP=INDEX VOLUME VTOC [,EXTENTS=addr] [,MAPRCDS=YES ¹ (YES,addr) NO ² (NO,addr)] [,UCB=(req) DEB=addr] [,COUNT=YES NO] [,CTAREA=addr] [,IOAREA=KEEP (KEEP,addr) NOKEEP (NOKEEP,addr)] [,BRANCH=NO YES ³ (YES,SUP) (YES,PGM)] [,MF=I L (E,addr)]
---------	---------	--

1 Default if MF=I.

2 Default if MF=L or MF=(E,addr).

3 Default is SUP if YES is coded.

ACCESS=MAPDATA: REQUEST INFORMATION FROM THE INDEX SPACE MAPS

ACCESS=MAPDATA

Obtains data from the index space maps. Three kinds of data are available:

- The number of format-0 DSCBs (the data is obtained from the VTOC map of DSCBs)
- The number of unallocated VIRs in the index (the data is obtained from the VTOC index map)
- The number (and location) of extents of unallocated pack space (the data is obtained from the VTOC pack space map)

MAP: IDENTIFY THE MAP TO BE ACCESSED

MAP=INDEX

Specifies that the VTOC index map (VIXM) is to be accessed and a count of unallocated VIRs returned. COUNT=YES must also be coded.

MAP=VOLUME

Specifies that the VTOC pack space map (VPSM) is to be accessed and information on unallocated extents of pack space returned. EXTENTS=addr and COUNT=NO must also be coded.

MAP=VTOC

Specifies that the VTOC map of DSCBs (VMDS) is to be accessed and a count of format-0 DSCBs returned. COUNT=YES must also be coded.

EXTENTS: IDENTIFY WHERE EXTENTS FROM THE VPSM ARE RETURNED

EXTENTS=addr

If one or more extents from the VPSM are requested, EXTENTS is the address of a 1-byte count field containing the number of 5-byte extents that follow. In the first two bytes of the first 5-byte extent, you must supply the relative track address (RTA) at which CVAF should start the VPSM search. The first extent area is updated with information on the next free extent found that has a higher starting RTA than that supplied. Each subsequent extent area is filled in with information on free space extents (in ascending track address order).

Information on free extents has the format, XXYYZ, where:

- XX is the relative track address of the first track of the extent.
- YY is the number of whole cylinders in the extent.
- Z is the number of additional tracks in the extent.

Only XX is supplied by the caller in the first extent area. CVAF will start searching the VPSM at relative track address XX.

If all the unallocated extents in the VPSM are provided before filling in all the supplied extent areas, the remaining extent areas are set to zero. Register 15 is set to 4 on return, with the CVSTAT field in the CVPL set to X'20' to indicate end of data.

MAPRCDS: KEEP OR FREE MAPRCDS BUFFER LIST AND BUFFERS

MAPRCDS=YES

Specifies that the buffer list and buffers are to be retained at the end of the function.

If YES is specified and no buffer list is supplied through the CVAF parameter list, CVAF will read the MAP VIRs into buffers obtained by CVAF. The buffer list that contains the address and RBAs of the VIRs can be accessed after the CVAF call from the CVAF parameter list field, CVMRCDS. The buffer list and VIR buffers are in the caller's protect key: subpool 0 if the caller is not authorized; subpool 229 if the caller is authorized.

YES is the default if MF=I is specified or defaulted.

When processing on the current volume is finished, release all areas that were kept.

MAPRCDS=(YES,addr)

If YES is coded, but the buffer list address (CVMRCDS in CVAF parameter list) is supplied, the VIRs are not read.

The CVMRCDS buffer list from one CVAF call can be passed to another CVAF macro call through the MAPRCDS keyword.

MAPRCDS=NO

If MAPRCDS=NO is coded, the MAP records buffers and buffer list will be freed upon completion of the CVAFDSM function.

NO is the default if MF=L is specified.

MAPRCDS=(NO,addr)

Causes buffer lists and buffers previously obtained by CVAF to be freed.

Buffer lists and buffers obtained by CVAF must be freed by the caller. This can be done in one of three ways:

- By coding MAPRCDS=NO on the call that obtained the buffers.
- By coding MAPRCDS=NO on a subsequent CVAF call.
- By calling CVAFDIR ACCESS=RLSE and providing the buffer list in the BUFLIST keyword.

If MF=(E,addr) is coded and MAPRCDS is not coded, the parameter list value of MAPRCDS is not changed.

Note: You must enqueue the VTOC and reserve the unit to maintain the integrity of the MAP records read.

UCB|DEB: SPECIFY THE VTOC TO BE ACCESSED

UCB=(reg)

Supplies the address of the UCB for the unit whose VTOC is to be accessed. An unauthorized caller may not supply a UCB to CVAF.

Note: Code the address of the UCB parameter only as register (2-12). Coding an RX-Type address here gives you unpredictable results.

DEB=addr

Specifies the address of a DEB opened to the VTOC you want to access. CVAF does not allow output requests to the VTOC or VTOC index if you specify the DEB subparameter. If you are not authorized, you cannot perform any asynchronous activity (such as EXCP, CLOSE, EOV) against the data set represented by the DEB because CVAF removes the DEB from the DEB table for the duration of the CVAF call. If you are not authorized (neither APF authorized nor in a system key), you must specify a DEB address, not a UCB, to CVAFDSM. See "Identifying the Volume" on page 26 for further details.

If you supply a previously obtained CVAF I/O area through the IOAREA keyword, you need not specify the UCB or DEB keyword. Otherwise you must specify either the UCB or DEB keyword. If you specify a UCB, the UCB address in the CVPL is overlaid by the UCB address in the I/O area.

If you supply both the UCB and DEB addresses in the CVPL, CVAF uses the DEB address and overlays the UCB address in the CVPL with the UCB address in the DEB.

COUNT: OBTAIN A COUNT OF UNALLOCATED DSCBS OR VIRS

COUNT=YES

Indicates that a count of unallocated DSCBs or VIRs in the designated space map is requested. MAP=VTOC or MAP=INDEX must be specified if COUNT=YES is coded.

COUNT=NO

Indicates that a count of unallocated DSCBs or VIRs is not desired but, rather, information on free space on the pack is desired. MAP=VOLUME must be coded if COUNT=NO is coded or defaulted.

CTAREA: SUPPLY A FIELD TO CONTAIN THE NUMBER OF FORMAT-0 DSCBS

CTAREA=addr

Gives the address of a 4-byte field to contain the number of format-0 DSCBs when COUNT=YES, MAP=VTOC is specified; or the number of unallocated VIRs in the VTOC index when COUNT=YES, MAP=INDEX is specified.

IOAREA: KEEP OR FREE THE I/O WORK AREA

IOAREA=KEEP

Specifies that the CVAF I/O area associated with the CVAF parameter list is to be kept upon completion of the CVAF request. IOAREA=KEEP may be coded with BRANCH=NO only if the caller is authorized (APF or system key).

If IOAREA=KEEP is coded, the caller must issue CVAF with IOAREA=NOKEEP specified at some future time, whether or not any further VTOC access is required: for example, the recovery routine of the caller of CVAF.

Coding IOAREA=KEEP allows subsequent CVAF requests to be more efficient, as certain initialization functions can be bypassed. Neither DEB nor UCB need be specified when a previously obtained CVAF I/O area is supplied; neither can they be changed.

When IOAREA=KEEP is first issued, CVAF returns the CVAF I/O area in the CVAF parameter list (CVIOAR). Subsequent calls of CVAF may use that same parameter list, and CVAF will obtain its I/O area from the CVIOAR.

When processing on the current volume is finished, release all areas that were kept.

IOAREA=(KEEP,addr)

Provides the address of a previously obtained I/O area. If a different CVAF parameter list is used, the previously obtained CVAF I/O area may be passed to CVAF by coding its address as the second parameter of the IOAREA keyword.

IOAREA=NOKEEP

Causes the work area to be freed upon completion of the CVAF request.

IOAREA=(NOKEEP,addr)

Causes a previously obtained work area to be freed upon completion of the CVAF request.

BRANCH: SPECIFY THE ENTRY TO THE MACRO

BRANCH=(YES,SUP)

Requests that the branch entry to CVAFDIR be used. The caller must be in supervisor state. Protect key checking is bypassed.

An 18-word save area must be supplied if BRANCH=YES is coded. No lock may be held on entry to CVAF. SRB mode is not allowed.

BRANCH=YES

Is equivalent to BRANCH=(YES,SUP), because SUP is the default when YES is coded. Protect key checking is bypassed.

BRANCH=(YES,PGM)

Requests the branch entry. The caller must be APF authorized and in problem state. Protect key checking is bypassed.

BRANCH=NO

Requests the SVC entry. The caller must be APF authorized if any output operations are requested. Protect key checking is performed.

MF: SPECIFY THE FORM OF THE MACRO

This keyword specifies whether the list, execute, or normal form of the macro is requested.

MF=I

If I is coded or if neither L nor E is coded, the CVAF parameter list is generated, as is code, to call CVAF. This is the normal form of the macro.

MF=L

L indicates the list form of the macro. A parameter list is generated, but code to call CVAF is not generated.

MF=(E,addr)

E indicates the execute form of the macro. The remote CVAF parameter list supplied as 'addr' is used in, and can be modified by, the execute form of the macro.

RETURN CODES FROM THE CVAFDSM MACRO

On return from CVAF, register 1 contains the address of the CVPL (CVAF parameter list), and register 15 contains one of the following return codes:

Code	Meaning
0(00)	The request was successful.
4(04)	End of data (CVSTAT is set to decimal 32), or an error was encountered. The CVSTAT field in the CVPL contains an indication of the cause of the error. (CVSTAT code descriptions are in Appendix C, "VTOC Index Error Message and Associated Codes" on page 221)
8(08)	Invalid VTOC index structure. CVSTAT contains an indication of the cause of the error. (CVSTAT code descriptions are in Appendix C, "VTOC Index Error Message and Associated Codes" on page 221)
12(0C)	The CVAF parameter list is not in your protect key or is invalid (the ID is invalid, or the length field is incorrect, or the CVFCTN field is invalid). The CVPL has not been modified.
16(10)	An I/O error was encountered.

CVAFSEQ MACRO

OVERVIEW OF THE CVAFSEQ MACRO

The CVAFSEQ macro may be used to:

- Read an indexed VTOC sequentially in data-set-name (DSN) order
- Read an indexed VTOC or a nonindexed VTOC in physical-sequential order

SYNTAX

<u>[label]</u>	CVAFSEQ	ACCESS=GT GTEQ [,BUFLIST= <u>addr</u>] [,DSN= <u>addr</u>] [,UCB=(<u>reg</u>) DEB= <u>addr</u>] [,DSNONLY=NO YES] [,ARG= <u>addr</u>] [,IOAREA=KEEP (KEEP, <u>addr</u>) NOKEEP (NOKEEP, <u>addr</u>)] [,IXRCDS=KEEP (KEEP, <u>addr</u>) NOKEEP (NOKEEP, <u>addr</u>)] [,BRANCH=NO YES ¹ (YES,SUP) (YES,PGM)] [,MF=I L (E, <u>addr</u>)]
----------------	---------	---

¹ If YES, default is SUP.

ACCESS: SPECIFY RELATIONSHIP BETWEEN SUPPLIED AND RETURNED DSN

ACCESS=GT

Specifies that the DSN or argument value is to be used to return a DSCB whose DSN or argument is greater than that supplied.

ACCESS=GTEQ

Specifies that the DSN or argument value is to be used to return a DSCB whose DSN or argument is greater than or equal to that supplied.

Note: A CVAF call specifying ACCESS=GTEQ should be followed by an ACCESS=GT request, or the same DSCB or name will be returned.

BUFLIST: SPECIFY ONE OR MORE BUFFER LISTS

BUFLIST=addr

The BUFLIST keyword supplies the address of a buffer list used to read or write DSCBs and VIRs.

DSN: SPECIFY ACCESS BY DSN ORDER OR BY PHYSICAL-SEQUENTIAL ORDER

DSN=addr

Specifies that access of an indexed VTOC is by DSN order. BUFLIST is required if DSNONLY=NO is coded or defaulted.

DSN omitted

If you omit the DSN keyword, access of an indexed or nonindexed VTOC is by physical-sequential order. BUFLIST is required.

Note: If the order is physical-sequential, you must initialize the argument field in the first buffer list entry to zero or to the argument of the DSCB. If the argument is zero (BFLEARG=00), the read begins at the start of the VTOC. You must be authorized

(APF or system key) to read multiple DSCBs with the CVAFSEQ macro. See "Initiating Physical-Sequential Access" on page 33 for more information.

UCB|DEB: SPECIFY THE VTOC TO BE ACCESSED

UCB=(reg)

Supplies the address of the UCB for the unit whose VTOC is to be accessed. An unauthorized caller may not supply a UCB to CVAF.

Note: Code the address of the UCB parameter only as register (2-12). Coding an RX-Type address here gives you unpredictable results.

DEB=addr

Specifies the address of a DEB opened to the VTOC you want to access. CVAF does not allow output requests to the VTOC or VTOC index if you specify the DEB subparameter. If you are not authorized, you cannot perform any asynchronous activity (such as EXCP, CLOSE, EOVS) against the data set represented by the DEB because CVAF removes the DEB from the DEB table for the duration of the CVAF call. If you are not authorized (neither APF authorized nor in a system key), you must specify a DEB address, not a UCB, to CVAFSEQ.

If you supply a previously obtained CVAF I/O area through the IOAREA keyword, you need not specify the UCB or DEB keyword. Otherwise you must specify either the UCB or DEB keyword. If you specify a UCB, the UCB address in the CVPL is overlaid by the UCB address in the I/O area.

If you supply both the UCB and DEB addresses in the CVPL, CVAF uses the DEB address and overlays the UCB address in the CVPL with the UCB address in the DEB.

DSNONLY: SPECIFY THAT ONLY THE DATA SET NAME BE READ

This keyword is applicable only to accessing an indexed VTOC in DSN order.

DSNONLY=NO

Requests that the data set name be obtained from the VTOC index and the DSCB be read into a buffer supplied through the BUFLIST keyword. BUFLIST is required.

DSNONLY=YES

Requests that only the data set name be obtained from the VTOC index. If the ARG keyword is coded, the argument of the DSCB is returned.

ARG: SPECIFY WHERE THE ARGUMENT OF THE DSCB IS TO BE RETURNED

This keyword is applicable only to accessing an indexed VTOC in DSN order with DSNONLY=YES coded.

ARG=addr

Provides the address of the 5-byte area where the CCHHR of each data set name in the VTOC index is returned when DSNONLY=YES is coded.

IOAREA: KEEP OR FREE THE I/O WORK AREA

IOAREA=KEEP

Specifies that the CVAF I/O area associated with the CVAF parameter list is to be kept upon completion of the CVAF request. IOAREA=KEEP may be coded with BRANCH=NO only if the caller is authorized (APF, or system key).

If IOAREA=KEEP is coded, the caller must issue CVAF with IOAREA=NOKEEP specified at some future time, whether or not any further VTOC access is required: for example, the recovery routine of the caller of CVAF.

Coding IOAREA=KEEP allows subsequent CVAF requests to be more efficient, because certain initialization functions can be bypassed. Neither DEB nor UCB need be specified when a previously obtained CVAF I/O area is supplied; neither can they be changed.

When IOAREA=KEEP is first issued, CVAF returns the CVAF I/O area in the CVAF parameter list (CVIOAR). Subsequent calls of CVAF may use that same parameter list, and CVAF will obtain its I/O area from the CVIOAR.

When processing on the current volume is finished, release all areas that were kept.

IOAREA=(KEEP,addr)

Provides the address of a previously obtained I/O area. If a different CVAF parameter list is used, the previously obtained CVAF I/O area may be passed to CVAF by coding its address as the second parameter of the IOAREA keyword.

IOAREA=NOKEEP

Causes the work area to be freed upon completion of the CVAF request.

IOAREA=(NOKEEP,addr)

Causes a previously obtained work area to be freed upon completion of the CVAF request.

IXRCDS: RETAIN VIERS IN VIRTUAL STORAGE

This keyword applies to an indexed VTOC only.

IXRCDS=KEEP

Specifies that the VIERS read into storage during the CVAF function are to be kept in virtual storage. The VIERS are retained even if the index function is unsuccessful. The VIERS are accessed from the CVAF parameter list (CVIRCDS). CVIRCDS is the address of a buffer list containing the VIR buffer addresses and RBAs of the VIERS read.

Index search function will dynamically update the buffer list and, when necessary, obtain additional buffer lists and chain them together.

If KEEP is specified and no buffer list is supplied to CVAF in the CVPL, CVAF will obtain a buffer list and buffers and read the first high-level VIER. The address of the buffer list is placed in the CVIRCDS field of the CVPL. The first high-level VIER will be checked for the VXFHLV bit and to see if the VXVISE bit is off.

The buffer list and VIR buffers are in the caller's protect key. The subpool is 0 if the caller is not authorized; subpool 229 if the caller is authorized.

If IXRCDS=KEEP for a nonindexed VTOC, a request to read a DSCB may be performed, but an error code will be returned.

When processing on the current volume is finished, release all areas that were kept.

IXRCDS=(KEEP,addr)

The CVIRCDS from one CVAF call can be passed to another CVAF parameter list by specifying the address as the second parameter in the IXRCDS keyword.

IXRCDS=NOKEEP

If NOKEEP is coded, the VIERS that are accessed (if any) are not retained. Furthermore, the buffer list supplied in the CVIRCDS field in the CVAF parameter list is released, as are all buffers found in the buffer list. If the skip bit is set in any entry in the buffer list, the buffer and buffer list will not be freed.

IXRCDS=(NOKEEP,addr)

Specifies that previously accessed VIERS are not to be retained.

You must free buffer lists and buffers obtained by CVAF. This can be done in one of three ways:

- By coding IXRCDS=NOKEEP on the CVAFSEQ macro that obtained the buffers
- By coding IXRCDS=NOKEEP on a subsequent CVAF macro
- By coding CVAFDIR ACCESS=RLSE and providing the address of the buffer list in the BUFLIST keyword

Note: You must enqueue the VTOC and reserve the unit to maintain the integrity of the VIERS read.

BRANCH: SPECIFY THE ENTRY TO THE MACRO

BRANCH=(YES,SUP)

Requests that the branch entry to CVAFDIR be used. The caller must be in supervisor state. Protect key checking is bypassed.

An 18-word save area must be supplied if BRANCH=YES is coded. No lock may be held on entry to CVAF. SRB mode is not allowed.

BRANCH=YES

Is equivalent to BRANCH=(YES,SUP), because SUP is the default when YES is coded. Protect key checking is bypassed.

BRANCH=(YES,PGM)

Requests the branch entry. The caller must be APF authorized and in problem state. Protect key checking is bypassed.

BRANCH=NO

Requests the SVC entry. The caller must be APF authorized if any output operations are requested. Protect key checking is performed.

MF: SPECIFY THE FORM OF THE MACRO

This keyword specifies whether the list, execute, or normal form of the macro is requested.

MF=I

If I is coded, or neither L nor E is coded, the CVAF parameter list is generated, as is code, to call CVAF. This is the normal form of the macro.

MF=L

L indicates the list form of the macro. A parameter list is generated, but code to call CVAF is not generated.

MF=(E,addr)

E indicates the execute form of the macro. The remote CVAF parameter list supplied as 'addr' is used in and can be modified by the execute form of the macro.

RETURN CODES FROM THE CVAFSEQ MACRO

On return from CVAF, register 1 contains the address of the CVPL (CVAF parameter list), and register 15 contains one of the following return codes:

Code	Meaning
00(X'00')	The request was successful.
04(X'04')	End of data (CVSTAT is set to decimal 32), or an error was encountered. The CVSTAT field in the CVPL contains an indication of the cause of the error. Error descriptions are in Appendix C, "VTOC Index Error Message and Associated Codes" on page 221
08(X'08')	Invalid VTOC index structure. CVSTAT contains an indication of the cause of the error. Error descriptions are in Appendix C, "VTOC Index Error Message and Associated Codes" on page 221
12(X'0C')	The CVPL (CVAF parameter list) is not in your protect key, or is invalid (the ID is invalid, or the length field is incorrect, or the CVFCTN field is invalid). The CVPL has not been modified.
16(X'10')	An I/O error was encountered.

CVAFTST MACRO

OVERVIEW OF THE CVAFTST MACRO

The CVAFTST macro determines whether the system supports an indexed VTOC, and, if it does, whether the VTOC on the unit whose UCB is supplied is indexed or nonindexed.

You will get a return code of 12 if CVAFTST cannot determine whether an indexed or nonindexed VTOC is on the unit's volume. You should not receive a return code of 12 from CVAFTST if you have opened a data set (including the VTOC) on the volume.

You need no authorization to issue the CVAFTST macro.

SYNTAX

<code>[[label]]</code>	<code>CVAFTST</code>	<code>UCB=(reg)</code>
------------------------	----------------------	------------------------

UCB: SPECIFY THE VTOC TO BE TESTED

`UCB=(reg)`

Supplies the address of the UCB for the volume whose VTOC is to be tested.

Note: Code the address of the UCB parameter only as register (2-12). Coding an RX-Type address here gives you unpredictable results.

RETURN CODES FROM THE CVAFTST MACRO

On return from CVAF, register 15 contains one of the following return codes:

Code	Meaning
0(00)	The system does not support an indexed VTOC. The volume should be considered to have a nonindexed VTOC. The UCB was not inspected to determine its validity or status.
4(04)	The system supports an indexed VTOC, but the volume has a nonindexed VTOC.
8(08)	The system supports an indexed VTOC and the volume has an indexed VTOC.
12(0C)	The system supports an indexed VTOC, but the volume is not mounted or the VIB is not initialized for it; thus, the status (indexed or nonindexed) of the VTOC cannot be determined.
16(10)	The system supports an indexed VTOC, but the unit is not a DASD or has a VIO UCB, or the UCB address is invalid.

| APPENDIX B. EXAMPLES OF VTOC ACCESS MACROS

The examples that follow are partial assembler listings that include expansions of each VTOC access macro. The expansions are provided to show how the VTOC macros can be substituted for existing procedures.

| EXAMPLE 1: USING THE CVAFDIR MACRO WITH AN INDEXED OR NON-INDEXED VTOC

This example uses the CVAFDIR macro to read a DSCB of a given data set name and determines whether the DSCB is for a partitioned data set. The address of the 44-byte data set name is supplied to the program in register 5 (labeled RDSN in the example). The address of a DEB open to the VTOC is supplied to the program in register 4 (labeled RDEB in the example).

The buffer list is in the program and is generated by the ICVAFBFL macro. The DSCB buffer is in the program and is generated by the IECSDSL1 macro.

```
EXAMPLE1 CSECT
        STM 14,12,12(RSAVE)
        BALR 12,0
        USING *,12
        ST  RSAVE,SAVEAREA+4
        LA  RWORK,SAVEAREA
        ST  RWORK,8(RSAVE)
        LR  RSAVE,RWORK
*****
*
*       REGISTERS
*
*****
REG1    EQU    1           REGISTER 1
RWORK  EQU    3           WORK REGISTER
RDEB   EQU    4           DEB ADDRESS
RDSN   EQU    5           ADDRESS OF DATA SET NAME
RSAVE  EQU    13          SAVE AREA ADDRESS
REG15  EQU    15          RETURN CODE REGISTER 15
*****
*
*       RETURN CODES
*
*****
PDSRTN EQU    0           DATA SET A PDS RETURN CODE
NOTFND  EQU    4           DATA SET NOT FOUND RETURN CODE
NOTPDS  EQU    8           DATA SET NOT A PDS RETURN CODE
UNEXPCD EQU    12          UNEXPECTED ERROR RETURN CODE
*****
*
*       READ DSCB INTO DSIFMTID.
*       DATA SET NAME ADDRESS SUPPLIED IN RDSN.
*       ADDRESS OF DEB OPEN TO VTOC SUPPLIED IN RDEB.
*       DETERMINE IF DATA SET IS A PARTITIONED DATA SET.
*       THIS PROGRAM IS NEITHER REENTRANT NOR REUSABLE.
*
*****
XC      BUFLIST(BFLHLN+BFELEN),BUFLIST ZERO BUFFER LIST
OI      BFLHFL,BFLHDSCB      DSCBS TO BE READ WITH BUFFER LIST
MVI     BFLHNOE,1           ONE BUFFER LIST ENTRY
LA      RWORK,DSIFMTID      ADDRESS OF DSCB BUFFER
```

```

ST   RWORK,BFLEBUF          PLACE IN BUFFER LIST
OI   BFLEFL,BFLECHR        CCHHR OF DSCB RETURNED BY CVAF
MVI  BFLELTH,DSCBLTH       DATA PORTION OF DSCB READ - DSN      *
                                     SUPPLIED IN CVPL
MVC  DS1DSNAM,0(RDSN)      MOVE IN DATA SET NAME TO WORKAREA
CVAFDIR ACCESS=READ,DSN=DS1DSNAM,BUFLIST=BUFLIST,DEB=(RDEB)
+   CNOP 0,4
+   BAL 1,ICV1E             LOAD CVPL LIST ADDRESS
+ICV1S EQU *               START OF CVPL
+   DC CL4'CVPL'           EBCDIC 'CVPL'
+   DC AL2(ICV1E-ICV1S)    LENGTH OF CVPL
+   DC XL1'01'             FUNCTION CODE
+   DC XL1'00'             STATUS INFORMATION
+   DC B'00000000'         FIRST FLAG BYTE
+   DC B'00000000'         SECOND FLAG BYTE
+   DC H'0'                RESERVED
+   DC A(0)                UCB ADDRESS
+   DC A(DS1DSNAM)         DATA SET NAME ADDRESS
+   DC A(BUFLIST)         BUFFER LIST ADDRESS
+   DC A(0)                INDEX VIR'S BUFFER LIST ADDRESS
+   DC A(0)                MAP VIR'S BUFFER LIST ADDRESS
+   DC A(0)                I/O AREA ADDRESS
+   DC A(0)                DEB ADDRESS
+   DC A(0)                ARGUMENT ADDRESS
+   DC A(0)                SPACE PARAMETER LIST ADDRESS
+   DC A(0)                EXTENT TABLE ADDRESS
+   DC A(0)                NEW VRF VIXM BUFFER LIST ADDR
+   DC A(0)                VRF DATA ADDRESS
+   DC A(0)                COUNT AREA ADDRESS
+ICV1E EQU *               END OF CVPL
+   ST RDEB,36(,1)        STORE DEB PTR IN PARM LIST
+   SVC 139
+   USING CVPL,REG1       ADDRESSABILITY TO CVPL
+   LTR REG15,REG15      ANY ERROR
+   BZ NOERROR           BRANCH IF NOT
*****
*
*   DETERMINE WHAT ERROR IS
*
*****
C   REG15,ERROR4         IS RETURN CODE 4
BNE OTHERERR           BRANCH IF NOT 4
CLI CVSTAT,STAT001     IS IT DATA SET NAME NOT FOUND?
BNE OTHERERR           BRANCH IF NOT
DROP REG1               ADDRESSABILITY TO CVPL NOT NEEDED
*****
*
*   DATA SET NAME NOT FOUND
*
*****
L   RSAVE,4(,RSAVE)
RETURN (14,12),RC=NOTFND SET UP DATA SET NOT FOUND ERROR
+   LM 14,12,12(13)      RESTORE THE REGISTERS
+   LA 15,NOTFND(0,0)    LOAD RETURN CODE
+   BR 14                RETURN
NOERROR EQU *           DSCB READ
MVC F1CCHHR,BFLEARG    MOVE CCHHR OF FORMAT 1/4 DSCB TO      *
                                     WORKAREA
CLI DS1FMTID,C'4'      IS DSCB A FORMAT 4 DSCB
BE NOTF1               BRANCH IF YES. NOT A FORMAT 1
TM DS1DSORG,DS1DSGPO   IS FORMAT 1 DSCB FOR PARTITIONED      *
                                     DATA SET
NOF1 BO PDS            BRANCH IF PDS
EQU *                 DSCB IS NOT A PDS
L   RSAVE,4(,RSAVE)
RETURN (14,12),RC=NOTPDS SET UP NOT PDS RETURN CODE
+   LM 14,12,12(13)      RESTORE THE REGISTERS

```

```

+ LA 15,NOTPDS(0,0) LOAD RETURN CODE
+ BR 14 RETURN
PDS EQU * DATA SET IS PARTITIONED
L RSAVE,4(,RSAVE)
RETURN (14,12),RC=PDSRTN SET UP PDS RETURN CODE
+ LM 14,12,12(13) RESTORE THE REGISTERS
+ LA 15,PDSRTN(0,0) LOAD RETURN CODE
+ BR 14 RETURN
OTHERERR EQU * UNEXPECTED ERROR
L RSAVE,4(,RSAVE)
RETURN (14,12),RC=UNEXPEC
+ LM 14,12,12(13) RESTORE THE REGISTERS
+ LA 15,UNEXPEC(0,0) LOAD RETURN CODE
+ BR 14 RETURN

```

```

ERROR4 DC F'4' ERROR RETURN CODE 4
BUFLIST ICVAFBFL DSECT=NO BUFFER LIST

```

```

+*****
+* BUFFER LIST HEADER
+*****

```

```

+BUFLIST DS OF BUFFER LIST HEADER
+BFLHNOE DS XL1 NUMBER OF ENTRIES
+BFLHFL DS XL1 KEY AND FLAG BYTE
+ ORG BFLHFL
+BFLHKEY DS XL1 PROTECT KEY (FIRST 4 BITS)
+BFLHVIR EQU X'08' BUF. LIST ENTRIES DESCRIBE VIRS
+BFLHDSCB EQU X'04' BUF. LIST ENTRIES DESCRIBE DSCBS
+ DS XL1 RESERVED
+BFLHSP DS XL1 SUBPOOL OF BUF. LIST/BUFFERS
+BFLHFCHN DS A FORWARD CHAIN PTR TO NEXT BUF.
+* LIST
+BFLHLN EQU *-BUFLIST LENGTH OF BUFFER LIST HEADER

```

```

+*****
+* BUFFER LIST ENTRY
+*****

```

```

+BFLE DS OF BUFFER LIST ENTRY
+BFLEFL DS XL1 BUFFER LIST ENTRY FLAG
+BFLERBA EQU X'80' ARGUMENT IS RBA
+BFLECHR EQU X'40' ARGUMENT IS CCHHR
+BFLETR EQU X'20' ARGUMENT IS TTR
+BFLEAUPD EQU X'10' CVAF UPDATED ARGUMENT FIELD
+BFLEMOD EQU X'08' DATA IN BUF. HAS BEEN MODIFIED
+BFLESKIP EQU X'04' SKIP THIS ENTRY
+BFLEIOER EQU X'02' I/O ERROR
+ DS XL1 RESERVED
+BFLELTH DS XL1 LENGTH OF DSCB BUFFER OR
+* LENGTH OF VIR DIVIDED BY 256
+BFLEARG DS XL5 ARGUMENT OF VIR OR DSCB (CCHHR)
+ ORG BFLEARG+1
+BFLATTR DS XL3 'TTR' OF ARGUMENT
+ ORG BFLEARG+1
+BFLEARBA DS XL4 'RBA' OF ARGUMENT
+BFLEBUF DS A BUFFER ADDRESS
+BFLELN EQU *-BFLE LENGTH OF A BUFFER LIST ENTRY
+IECSDSL1 (1) FORMAT 1 DSCB DATA SET NAME AND *
+* BUFFER

```

```

+IECSDSL1 EQU *
+IECSDSF1 EQU IECSDSL1
+DS1DSNAM DS CL44 DATA SET NAME
+DS1FMTID DS CL1 FORMAT IDENTIFIER
+DS1DSSN DS CL6 DATA SET SERIAL NUMBER
+DS1VOLSQ DS XL2 VOLUME SEQUENCE NUMBER
+DS1CREDT DS XL3 CREATION DATE
+DS1EXPDT DS XL3 EXPIRATION DATE
+DS1NOEPV DS XL1 NUMBER OF EXTENTS ON VOLUME
+DS1NOBDB DS XL1 NUMBER OF BYTES USED IN LAST
+* DIRECTORY BLOCK

```

+	DS	XL1	RESERVED	
+DS1SYSCD	DS	CL13	SYSTEM CODE	
+	DS	XL7	RESERVED	
+DS1DSORG	DS	XL2	DATA SET ORGANIZATION	
+*			DS1DSORG	
+DS1DSGIS	EQU	X'80'	IS - INDEXED SEQUENTIAL ORGANIZATION	@01A
+*				
+DS1DSGPS	EQU	X'40'	PS - PHYSICAL SEQUENTIAL ORGANIZATION	@01A
+*				
+DS1DSGDA	EQU	X'20'	DA - DIRECT ORGANIZATION	@01A
+DS1DSGCX	EQU	X'10'	CX - BTAM OR QTAM LINE GROUP	@01A
+*	EQU	X'08'	RESERVED	@01A
+*	EQU	X'04'	RESERVED	@01A
+DS1DSGPO	EQU	X'02'	PO - PARTITIONED ORGANIZATION	@01A
+DS1DSGU	EQU	X'01'	U - UNMOVABLE, THE DATA CONTAINS LOCATION DEPENDENT INFORMATION	@01A
+*				
+*				
+*				
			SECOND BYTE OF DS1DSORG	
+DS1DSGGS	EQU	X'80'	GS - GRAPHICS ORGANIZATION	@01A
+DS1DSGTX	EQU	X'40'	TX - TCAM LINE GROUP	@01A
+DS1DSGTQ	EQU	X'20'	TQ - TCAM MESSAGE QUEUE	@01A
+*	EQU	X'10'	RESERVED	@01A
+DS1ACBM	EQU	X'08'	ACCESS METHOD CONTROL BLOCK	@01A
+DS1DSGTR	EQU	X'04'	TR - TCAM 3705	@01A
+*	EQU	X'02'	RESERVED	@01A
+*	EQU	X'01'	RESERVED	@01A
+DS1RECFM	DS	XL1	RECORD FORMAT	
+DS1OPTCD	DS	XL1	OPTION CODE	
+DS1BLKL	DS	XL2	BLOCK LENGTH	
+DS1LRECL	DS	XL2	RECORD LENGTH	
+DS1KEYL	DS	XL1	KEY LENGTH	
+DS1RKP	DS	XL2	RELATIVE KEY POSITION	
+DS1DSIND	DS	XL1	DATA SET INDICATORS	
+DS1SCALO	DS	XL4	SECONDARY ALLOCATION	
+DS1LSTAR	DS	XL3	LAST USED TRACK AND BLOCK ON TRACK	
+DS1TRBAL	DS	XL2	BYTES REMAINING ON LAST TRACK USED	
+	DS	XL2	RESERVED	
+DS1EXT1	DS	XL10	FIRST EXTENT DESCRIPTION	
+*			EXTENT TYPE INDICATOR	
+*			EXTENT SEQUENCE NUMBER	
+*			LOWER LIMIT	
+*			UPPER LIMIT	
+DS1EXT2	DS	XL10	SECOND EXTENT DESCRIPTION	
+DS1EXT3	DS	XL10	THIRD EXTENT DESCRIPTION	
+DS1PTRDS	DS	XL5	POSSIBLE PTR TO A FORMAT 2 OR 3 DSCB	
+DS1END	EQU	*		
DSCBLTH	EQU	*-IECSDSL1-L'DS1DSNAM	LENGTH OF DATA PORTION OF DSCB	
FLCCHR	DS	XL5	CCHHR OF DSCB	
SAVEAREA	DS	18F	SAVE AREA	
CVPL	ICVAFPL	,	CVPL MAPPING MACRO	

+*****
+* CVAF PARAMETER LIST
+*****

+CVPL	DSECT		CVAF PARAMETER LIST
+	DS	0F	
+CVLBL	DS	CL4	EBCDIC 'CVPL'
+CVLTH	DS	H	LENGTH OF CVPL
+CVFCTN	DS	XL1	FUNCTION BYTE
+CVDIRD	EQU	X'01'	CVAFDIR ACCESS=READ
+CVDIWR	EQU	X'02'	CVAFDIR ACCESS=WRITE
+CVDIRLS	EQU	X'03'	CVAFDIR ACCESS=RLSE
+CVSEQGT	EQU	X'04'	CVAFSEQ ACCESS=GT
+CVSEQGTE	EQU	X'05'	CVAFSEQ ACCESS=GTEQ
+CVDMIXA	EQU	X'06'	CVAFDSM ACCESS=IXADD
+CVDMIXD	EQU	X'07'	CVAFDSM ACCESS=IXDLT
+CVDMALC	EQU	X'08'	CVAFDSM ACCESS=ALLOC
+CVDMRLS	EQU	X'09'	CVAFDSM ACCESS=RLSE
+CVDMMAP	EQU	X'0A'	CVAFDSM ACCESS=MAPDATA
+CVVOL	EQU	X'0B'	CVAFVOL ACCESS=VIBBLD

+CVVFRD	EQU	X'0C'	CVAFVRF ACCESS=READ
+CVVRFWR	EQU	X'0D'	CVAFVRF ACCESS=WRITE
+CVSTAT	DS	XL1	STATUS INFORMATION (SEE LIST * BELOW)
+			FIRST FLAG BYTE
+CVFL1	DS	XL1	INDEXED VTOC ACCESSED
+CV1IVT	EQU	X'80'	IOAREA=KEEP
+CV1IOAR	EQU	X'40'	BRANCH=(YES,PGM)
+CV1PGM	EQU	X'20'	MAPRCDS=YES
+CV1MRCDS	EQU	X'10'	IXRCDS=KEEP
+CV1IRCDS	EQU	X'08'	MAP=INDEX
+CV1MAPIX	EQU	X'04'	MAP=VTOC
+CV1MAPVT	EQU	X'02'	MAP=VOLUME
+CV1MAPVL	EQU	X'01'	SECOND FLAG BYTE
+CVFL2	DS	XL1	HIVIER=YES
+CV2HIVIE	EQU	X'80'	VRF DATA EXISTS
+CV2VRF	EQU	X'40'	COUNT=YES
+CV2CNT	EQU	X'20'	RECOVER=YES
+CV2RCVR	EQU	X'10'	SEARCH=YES
+CV2SRCH	EQU	X'08'	DSNONLY=YES
+CV2DSNLY	EQU	X'04'	VERIFY=YES
+CV2VER	EQU	X'02'	OUTPUT-NEW HIGHEST LEVEL VIER
+CV2NLEVL	EQU	X'01'	CREATED
+*			RESERVED
+	DS	H	UCB ADDRESS
+CVUCB	DS	A	DATA SET NAME ADDRESS
+CVDSN	DS	A	BUFFER LIST ADDRESS
+CVBUFL	DS	A	INDEX VIR'S BUFFER LIST ADDRESS
+CVIRCDS	DS	A	MAP VIR'S BUFFER LIST ADDRESS
+CVMRCDS	DS	A	I/O AREA ADDRESS
+CVIOAR	DS	A	DEB ADDRESS
+CVDEB	DS	A	ARGUMENT ADDRESS
+CVARG	DS	A	SPACE PARAMETER LIST ADDRESS
+CVSPACE	DS	A	EXTENT TABLE ADDRESS
+CVEXTS	DS	A	NEW VRF VIXM BUFFER LIST ADDR
+CVBUFL2	DS	A	VRF DATA ADDRESS
+CVVRFDA	DS	A	COUNT AREA ADDRESS
+CVCTAR	DS	A	
+CVPLNGTH	EQU	*-CVPL	

+* VALUES OF CVSTAT
+*(THIS PART OF THE ICVAFPL MACRO EXPANSION IS NOT SHOWN)
END

EXAMPLE 2: USING THE CVAFDIR MACRO WITH AN INDEXED VTOC

This example uses the CVAFDIR macro to read one or more DSCBs on a VTOC. The UCB is supplied to the program in register 4 (labeled RUCB). The TTR of each DSCB read is to be returned to the caller. This program must be APF authorized.

The address of a parameter list is supplied to the program in register 5 (labeled RLIST). The parameter list contains one or more 3-word entries. The format of each 3-word entry is mapped by the LISTMAP DSECT. The first word contains the address of the data set name of the DSCB to be read. The second word contains the address of the 96-byte buffer into which the DSCB is to be read. The third word contains the address of the 3-byte TTR of the DSCB read.

The CVPL is generated by a list form of the CVAFDIR macro at label CVPL. The BUFLIST, IXRCDS, IOAREA, and BRANCH keywords are coded on the list form of the macro. IXRCDS=KEEP and IOAREA=KEEP are coded to avoid overhead if two or more DSCBs are to be read. BRANCH=(YES,PGM) is coded in the list form of the CVAFDIR macro to cause the CVPL to have the CV1PGM bit set to one; this will indicate to CVAF that the caller is authorized by APF and not in supervisor state. The execute forms of the CVAFDIR macro then specify BRANCH=YES, and not BRANCH=(YES,PGM), because the CV1PGM bit is set in the list form of the macro.

The CVAFDIR macro with ACCESS=RLSE is coded before the program exits in order to release the CVAF I/O area and the index records buffer list. BUFLIST=0 is coded because no user-supplied buffer list is to be released; BUFLIST was coded on the list form of the CVAFDIR macro and, therefore, is in the CVBUFL field of the CVPL. This field must be set to zero for the release.

EXAMPLE2 CSECT

```

STM 14,12,12(13)
BALR 12,0
USING *,12
ST 13,SAVEAREA+4
LA RWORK,SAVEAREA
ST RWORK,8(,13)
LR 13,RWORK

```

*
* REGISTERS
*

```

RWORK EQU 3                   WORK REGISTER
RUCB EQU 4                   UCB ADDRESS SUPPLIED BY CALLER
RLIST EQU 5                  ADDRESS OF PARAMETER LIST
RDSN EQU 6                  ADDRESS OF DATA SET NAME
RTTR EQU 7                  ADDRESS OF TTR
REG15 EQU 15                 RETURN CODE REGISTER 15

```

*
* READ DSCB OF DATA SET NAME SUPPLIED. RETURN TTR OF DSCB.
* UCB ADDRESS SUPPLIED IN RUCB.
* ADDRESS OF PARAMETER LIST IN RLIST.
* WORD 1 OF PARAMETER LIST = ADDRESS OF DATA SET NAME
* WORD 2 OF PARAMETER LIST = ADDRESS OF DSCB TO BE RETURNED
* WORD 3 OF PARAMETER LIST = ADDRESS OF TTR TO BE RETURNED
* WORDS 1-3 DUPLICATED WITH THE HIGH ORDER BIT OF
* WORD 3 SET TO ONE FOR LAST ENTRY.
*

 USING LISTMAP,RLIST ADDRESSABILITY TO PARMLIST

```

TOPLOOP EQU * LOOP FOR EACH DSCB
XC BUFLIST(BFLHLN+BFILELN),BUFLIST ZERO BUFFER LIST
OI BFLHFL,BFLHDSCB DSCBS TO BE READ WITH BUFFER LIST
MVI BFLHNOE,1 ONE BUFFER LIST ENTRY
LA RWORK,LISTDSCB ADDRESS OF DSCB BUFFER
ST RWORK,BFLEBUF PLACE IN BUFFER LIST
OI BFLEFL,BFLETTR TTR OF DSCB RETURNED BY CVAF
MVI BFLELTH,DSCBLTH DATA PORTION OF DSCB READ - DSN *
                          SUPPLIED IN CVPL
L RDSN,LISTDSN ADDRESS OF DATA SET NAME
CVAFDIR DSN=(RDSN),UCB=(RUCB),MF=(E,CVPL),BRANCH=YES
+ LA 1,CVPL LOAD PARAMETER REG 1
+ ST RUCB,12(,1) STORE UCB PTR IN PARM LIST
+ ST RDSN,16(,1) STORE DSN PTR IN PARM LIST
+ L 15,16 LOAD THE CVT
+ L 15,328(,15) LOAD VS1/VS2 COMMON EXTENSION2
+ L 15,12(,15) LOAD THE CVT CVAF TABLE
+ L 15,0(,15) LOAD THE CVAF ADDRESS
+ BALR 14,15 BRANCH AND LINK TO CVAF
L RTTR,LISTTTR ADDRESS OF TTR TO BE RETURNED
USING TTRMAP,RTTR MAP OF TTR
LTR REG15,REG15 ANY ERROR
BZ NOERROR BRANCH IF NOT
XC TTR,TTR ZERO TTR INDICATING NO DSCB
B RELOOP GET NEXT ENTRY
NOERROR EQU * DSCB READ
MVC TTR,BFLEARG RETURN TTR OF DSCB
RELOOP EQU * GET NEXT ENTRY
TM LASTLIST,LASTBIT IS IT LAST ENTRY IN LIST?
LA RLIST,NEXTLIST GET NEXT ENTRY
BZ TOPLOOP PROCESS NEXT LIST
CVAFDIR ACCESS=RLSE, RELEASE CVAF OBTAINED AREAS *
IOAREA=NOKEEP, RELEASE IOAREA *
IXRCDS=NOKEEP, RELEASE VIER BUFFER LIST *
BUFLIST=0, NO USER BUFFER LIST SUPPLIED TO RLSE*
BRANCH=YES, BRANCH ENTER CVAF *
MF=(E,CVPL)
+ LA 1,CVPL LOAD PARAMETER REG 1
+ MVI 6(1),X'03' SET FUNCTION CODE
+ NI 8(1),B'10110111' RESET CVAF FLAGS OFF
+ LA 15,0 GET BUFLIST ADDRESS AND
+ ST 15,20(,1) STORE BUFLIST PTR IN PARM LIST
+ L 15,16 LOAD THE CVT
+ L 15,328(,15) LOAD VS1/VS2 COMMON EXTENSION2
+ L 15,12(,15) LOAD THE CVT CVAF TABLE
+ L 15,0(,15) LOAD THE CVAF ADDRESS
+ BALR 14,15 BRANCH AND LINK TO CVAF
L 13,SAVEAREA+4
RETURN (14,12)
+ LM 14,12,12(13) RESTORE THE REGISTERS
+ BR 14 RETURN

```

BUFLIST ICVAFBFL DSECT=NO BUFFER LIST

```

+*****
+* BUFFER LIST HEADER
+*****

```

```

+BUFLIST DS OF BUFFER LIST HEADER
+BFLHNOE DS XL1 NUMBER OF ENTRIES
+BFLHFL DS XL1 KEY AND FLAG BYTE
+ ORG BFLHFL
+BFLHKEY DS XL1 PROTECT KEY (FIRST 4 BITS)
+BFLHVIR EQU X'08' BUF. LIST ENTRIES DESCRIBE VIRS
+BFLHDSCB EQU X'04' BUF. LIST ENTRIES DESCRIBE DSCBS
+ DS XL1 RESERVED
+BFLHSP DS XL1 SUBPOOL OF BUF. LIST/BUFFERS
+BFLHFCHN DS A FORWARD CHAIN PTR TO NEXT BUF.
+* LIST
+BFLHLN EQU *-BUFLIST LENGTH OF BUFFER LIST HEADER

```

```

+*****
+*          BUFFER LIST ENTRY
+*****

+BFLE      DS      0F          BUFFER LIST ENTRY
+BFLEFL    DS      XL1        BUFFER LIST ENTRY FLAG
+BFLELBA   EQU     X'80'      ARGUMENT IS RBA
+BFLECHR   EQU     X'40'      ARGUMENT IS CCHHR
+BFLETR    EQU     X'20'      ARGUMENT IS TTR
+BFLEAUPD  EQU     X'10'      CVAF UPDATED ARGUMENT FIELD
+BFLEMOD   EQU     X'08'      DATA IN BUF. HAS BEEN MODIFIED
+BFLESKIP  EQU     X'04'      SKIP THIS ENTRY
+BFLEIOER  EQU     X'02'      I/O ERROR
+          DS      XL1        RESERVED
+BFLELTH   DS      XL1        LENGTH OF DSCB BUFFER OR
+*          DS      XL1        LENGTH OF VIR DIVIDED BY 256
+BFLEARG   DS      XL5        ARGUMENT OF VIR OR DSCB (CCHHR)
+          ORG     BFLEARG+1
+BFLEATTR  DS      XL3        'TTR' OF ARGUMENT
+          ORG     BFLEARG+1
+BFLEARBA  DS      XL4        'RBA' OF ARGUMENT
+BFLEBUF   DS      A          BUFFER ADDRESS
+BFLELN    EQU     *-BFLE     LENGTH OF A BUFFER LIST ENTRY
+SAVEAREA  DS      18F       REGISTER SAVE AREA
+LISTMAP   DSECT
+LISTDSN   DS      F          ADDRESS OF DATA SET NAME
+LISTDSCB  DS      F          ADDRESS OF BUFFER FOR DSCB TO BE *
+          DS      F          RETURNED
+LISTTTR   DS      0F        ADDRESS OF TTR OF DSCB TO BE *
+          DS      0F        RETURNED
+LASTLIST  DS      X          FIRST BYTE
+LASTBIT   EQU     X'80'      LAST ENTRY IN LIST
+          DS      XL3        REMAINDER OF TTR ADDRESS
+NEXTLIST  EQU     *          NEXT LIST
+DSCB      DSECT
+IECSDSL1  EQU     *          IECSDSL1 (1)
+          EQU     *          FORMAT 1 DSCB
+IECSDSF1  EQU     IECSDSL1
+DS1DSNAM  DS      CL44      DATA SET NAME
+DS1FMTID  DS      CL1       FORMAT IDENTIFIER
+DS1DSSN   DS      CL6       DATA SET SERIAL NUMBER
+DS1VOLSQ  DS      XL2       VOLUME SEQUENCE NUMBER
+DS1CREDT  DS      XL3       CREATION DATE
+DS1EXPDT  DS      XL3       EXPIRATION DATE
+DS1NOEPV  DS      XL1       NUMBER OF EXTENTS ON VOLUME
+DS1NOBDB  DS      XL1       NUMBER OF BYTES USED IN LAST
+*          DS      XL1       DIRECTORY BLOCK
+          DS      XL1       RESERVED
+DS1SYSCD  DS      CL13      SYSTEM CODE
+          DS      XL7       RESERVED
+DS1DSORG  DS      XL2       DATA SET ORGANIZATION
+*          DS      XL2       FIRST BYTE OF DS1DSORG
+DS1DSGIS  EQU     X'80'     IS - INDEXED SEQUENTIAL          @01A
+*          EQU     X'80'     ORGANIZATION
+DS1DSGPS  EQU     X'40'     PS - PHYSICAL SEQUENTIAL          @01A
+*          EQU     X'40'     ORGANIZATION
+DS1DSGDA  EQU     X'20'     DA - DIRECT ORGANIZATION          @01A
+DS1DSGCX  EQU     X'10'     CX - BTAM OR QTAM LINE GROUP      @01A
+*          EQU     X'08'     RESERVED                          @01A
+*          EQU     X'04'     RESERVED                          @01A
+DS1DSGPO  EQU     X'02'     PO - PARTITIONED ORGANIZATION      @01A
+DS1DSGU   EQU     X'01'     U - UNMOVABLE, THE DATA          @01A
+*          EQU     X'01'     CONTAINS LOCATION DEPENDENT
+*          EQU     X'01'     INFORMATION
+*
+*
+*

```

```

+*                SECOND BYTE OF DSIDSORG
+DS1DSGGS EQU X'80' GS - GRAPHICS ORGANIZATION @01A
+DS1DSGTX EQU X'40' TX - TCAM LINE GROUP @01A
+DS1DSGTQ EQU X'20' TQ - TCAM MESSAGE QUEUE @01A
+* EQU X'10' RESERVED @01A
+DS1ACBM EQU X'08' ACCESS METHOD CONTROL BLOCK @01A
+DS1DSGTR EQU X'04' TR - TCAM 3705 @01A
+* EQU X'02' RESERVED @01A
+* EQU X'01' RESERVED @01A
+DS1RECFM DS XL1 RECORD FORMAT
+DS1OPTCD DS XL1 OPTION CODE
+DS1BLKL DS XL2 BLOCK LENGTH
+DS1LRECL DS XL2 RECORD LENGTH
+DS1KEYL DS XL1 KEY LENGTH
+DS1RKP DS XL2 RELATIVE KEY POSITION
+DS1DSIND DS XL1 DATA SET INDICATORS
+DS1SCALO DS XL4 SECONDARY ALLOCATION
+DS1LSTAR DS XL3 LAST USED TRACK AND BLOCK ON TRACK
+DS1TRBAL DS XL2 BYTES REMAINING ON LAST TRACK USED
+ DS XL2 RESERVED
+DS1EXT1 DS XL10 FIRST EXTENT DESCRIPTION
+* FIRST BYTE EXTENT TYPE INDICATOR
+* SECOND BYTE EXTENT SEQUENCE NUMBER
+* THIRD - SIXTH BYTES LOWER LIMIT
+* SEVENTH - TENTH BYTES UPPER LIMIT
+DS1EXT2 DS XL10 SECOND EXTENT DESCRIPTION
+DS1EXT3 DS XL10 THIRD EXTENT DESCRIPTION
+DS1PTRDS DS XL5 POSSIBLE PTR TO A FORMAT 2 OR 3 DSCB
+DS1END EQU *
DSCBLTH EQU *-DSCB-L'DS1DSNAM LENGTH OF DATA PORTION OF DSCB
TTRMAP DSECT
TTR DS XL3 TTR TO BE RETURNED
EXAMPLE2 CSECT
CVPL CVAFDIR ACCESS=READ,BUFLIST=BUFLIST,MF=L, *
      IOAREA=KEEP, KEEP IOAREA TO AVOID OVERHEAD *
      IXRCDS=KEEP CALLED IN PROGRAM STATE BUT APF *
      AUTHORIZED SO UCB IS SUPPLIED *

```

```

+ CNOP 0,4
+CVPL EQU *
+ DC CL4'CVPL' EBCDIC 'CVPL'
+ DC AL2(ICV8E-CVPL) LENGTH OF CVPL
+ DC XL1'01' FUNCTION CODE
+ DC XL1'00' STATUS INFORMATION
+ DC B'01001000' FIRST FLAG BYTE
+ DC B'00000000' SECOND FLAG BYTE
+ DC H'0' RESERVED
+ DC A(0) UCB ADDRESS
+ DC A(0) DATA SET NAME ADDRESS
+ DC A(BUFLIST) BUFFER LIST ADDRESS
+ DC A(0) INDEX VIR'S BUFFER LIST ADDRESS
+ DC A(0) MAP VIR'S BUFFER LIST ADDRESS
+ DC A(0) I/O AREA ADDRESS
+ DC A(0) DEB ADDRESS
+ DC A(0) ARGUMENT ADDRESS
+ DC A(0) SPACE PARAMETER LIST ADDRESS
+ DC A(0) EXTENT TABLE ADDRESS
+ DC A(0) NEW VRF VIXM BUFFER LIST ADDR
+ DC A(0) VRF DATA ADDRESS
+ DC A(0) COUNT AREA ADDRESS
+ICV8E EQU *
      ORG CVPL OVERLAY CVPL WITH EXPANSION OF MAP
CVPLMAP ICVAFPL DSECT=NO
+*****
+* CVAF PARAMETER LIST
+*****

```

```

+CVPLMAP DS OF CVAF PARAMETER LIST
+CVLBL DS CL4 EBCDIC 'CVPL'
+CVLTH DS H LENGTH OF CVPL

```

```

+CVFCTN DS XL1
+CVDIRD EQU X'01'
+CVDIWR EQU X'02'
+CVDIRLS EQU X'03'
+CVSEQGT EQU X'04'
+CVSEQGT EQU X'05'
+CVDMIXA EQU X'06'
+CVDMIXD EQU X'07'
+CVDMALC EQU X'08'
+CVDMRLS EQU X'09'
+CVDMMAP EQU X'0A'
+CVVOL EQU X'0B'
+CVVFRFD EQU X'0C'
+CVVRFWR EQU X'0D'
+CVSTAT DS XL1
+
+CVFL1 DS XL1
+CV1IVT EQU X'80'
+CV1IOAR EQU X'40'
+CV1PGM EQU X'20'
+CV1MRCDS EQU X'10'
+CV1IRCDS EQU X'08'
+CV1MAPIX EQU X'04'
+CV1MAPVT EQU X'02'
+CV1MAPVL EQU X'01'
+CVFL2 DS XL1
+CV2HIVIE EQU X'80'
+CV2VRF EQU X'40'
+CV2CNT EQU X'20'
+CV2RCVR EQU X'10'
+CV2SRCH EQU X'08'
+CV2DSNLY EQU X'04'
+CV2VER EQU X'02'
+CV2NLEVL EQU X'01'
+*
+ DS H
+CVUCB DS A
+CVDSN DS A
+CVBUFL DS A
+CVIRCDS DS A
+CVMRCDS DS A
+CVIOAR DS A
+CVDEB DS A
+CVARG DS A
+CVSPACE DS A
+CVEXTS DS A
+CVBUFL2 DS A
+CVVRFDA DS A
+CVCTAR DS A
+CVPLNGTH EQU *-CVPLMAP

```

```

+* VALUES OF CVSTAT
+*(THIS PART OF THE ICVAFPL MACRO EXPANSION IS NOT SHOWN)
END

```

```

FUNCTION BYTE
CVAFDIR ACCESS=READ
CVAFDIR ACCESS=WRITE
CVAFDIR ACCESS=RLSE
CVAFSEQ ACCESS=GT
CVAFSEQ ACCESS=GTEQ
CVAFDSM ACCESS=IXADD
CVAFDSM ACCESS=IXDLT
CVAFDSM ACCESS=ALLOC
CVAFDSM ACCESS=RLSE
CVAFDSM ACCESS=MAPDATA
CVAFVOL ACCESS=VIBBLD
CVAFVRF ACCESS=READ
CVAFVRF ACCESS=WRITE
STATUS INFORMATION (SEE LIST *)
BELOW)
FIRST FLAG BYTE
INDEXED VTOC ACCESSED
IOAREA=KEEP
BRANCH=(YES,PGM)
MAPRCDS=YES
IXRCDS=KEEP
MAP=INDEX
MAP=VTOC
MAP=VOLUME
SECOND FLAG BYTE
HIVIER=YES
VRF DATA EXISTS
COUNT=YES
RECOVER=YES
SEARCH=YES
DSNONLY=YES
VERIFY=YES
OUTPUT-NEW HIGHEST LEVEL VIER
CREATED
RESERVED
UCB ADDRESS
DATA SET NAME ADDRESS
BUFFER LIST ADDRESS
INDEX VIR'S BUFFER LIST ADDRESS
MAP VIR'S BUFFER LIST ADDRESS
I/O AREA ADDRESS
DEB ADDRESS
ARGUMENT ADDRESS
SPACE PARAMETER LIST ADDRESS
EXTENT TABLE ADDRESS
NEW VRF VIXM BUFFER LIST ADDR
VRF DATA ADDRESS
COUNT AREA ADDRESS

```

EXAMPLE 3: USING THE CVAFSEQ MACRO WITH AN INDEXED VTOC

This example uses the CVAFSEQ macro to count the number of ISAM data sets whose data set names are within the range defined by two supplied data set names. The addresses of the two data set names are supplied to the program in registers 6 and 7, labeled RDSN1 and RDSN2, respectively. The address of a DEB open to the VTOC is supplied in register 4, labeled RDEB.

The CVAF parameter list is expanded by a list form of the CVAFSEQ macro. ACCESS=GTEQ is specified on the list form of the macro and is, therefore, not coded in the first execution of the CVPL. Subsequent executions of the CVPL (at label RELOOP) specify ACCESS=GT.

End of data is tested by comparing the CVSTAT field to the value of STAT032, which is an equate in the ICVAFPL mapping macro.

The count of ISAM DSCBs matching the data set name criterion is returned in register 15, unless an error is encountered, in which case a negative 1 is returned in register 15.

```
EXAMPLE3 CSECT
          STM 14,12,12(13)
          BALR 12,0
          USING *,12
          ST 13,SAVEAREA+4
          LA RWORK,SAVEAREA
          ST RWORK,8(,13)
          LR 13,RWORK
```

XX

```
*
*          REGISTERS
*
```

XX

```
REG1 EQU 1          REGISTER 1
RWORK EQU 3        WORK REGISTER
RDEB EQU 4         DEB ADDRESS
RDSN1 EQU 6        ADDRESS OF DATA SET NAME 1
RDSN2 EQU 7        ADDRESS OF DATA SET NAME 2
REG15 EQU 15       RETURN CODE REGISTER 15
```

XX

```
*
*          COUNT THE NUMBER OF ISAM DATA SETS WHOSE DATA SET NAMES ARE
*          BETWEEN DSN1 AND DSN2 INCLUSIVELY.
*          RDSN1 CONTAINS ADDRESS OF DSN1.
*          RDSN2 CONTAINS ADDRESS OF DSN2.
*          ADDRESS OF DEB OPEN TO VTOC SUPPLIED IN RDEB.
*
```

XX

```
XC BUFLIST(BFLHLN+BFLELN),BUFLIST ZERO BUFFER LIST
OI BFLHFL,BFLHDSCB    DSCBS TO BE READ WITH BUFFER LIST
MVI BFLHNOE,1         ONE BUFFER LIST ENTRY
LA RWORK,DS1FMTID     ADDRESS OF DSCB BUFFER
ST RWORK,BFLEBUF      PLACE IN BUFFER LIST
MVI BFLELTH,DSCBLTH   DATA PORTION OF DSCB READ - DSN *
                          SUPPLIED IN CVPL
MVC DS1DSNAM,0(RDSN1) MOVE IN STARTING DATA SET NAME TO *
                          WORKAREA
XR RWORK,RWORK        ZERO COUNT
CVAFSEQ DEB=(RDEB),   FIND FIRST DATA SET WHOSE DATA SET *
          BUFLIST=BUFLIST, NAME IS GREATER THAN OR EQUAL TO *
          MF=(E,CVPL)  THAT OF DSN1
+ LA 1,CVPL           LOAD PARAMETER REG 1
+ ST RDEB,36(,1)     STORE DEB PTR IN PARM LIST
+ SVC 139
```

 +* BUFFER LIST HEADER

+BUFLIST	DS	OF	BUFFER LIST HEADER
+BFLHNOE	DS	XL1	NUMBER OF ENTRIES
+BFLHFL	DS	XL1	KEY AND FLAG BYTE
+	ORG	BFLHFL	
+BFLHKEY	DS	XL1	PROTECT KEY (FIRST 4 BITS)
+BFLHVIR	EQU	X'08'	BUF. LIST ENTRIES DESCRIBE VIRS
+BFLHDSCB	EQU	X'04'	BUF. LIST ENTRIES DESCRIBE DSCBS
+	DS	XL1	RESERVED
+BFLHSP	DS	XL1	SUBPOOL OF BUF. LIST/BUFFERS
+BFLHFCHN	DS	A	FORWARD CHAIN PTR TO NEXT BUF.
+*			LIST
+BFLHLN	EQU	*-BUFLIST	LENGTH OF BUFFER LIST HEADER

 +* BUFFER LIST ENTRY

+BFLE	DS	OF	BUFFER LIST ENTRY
+BFLEFL	DS	XL1	BUFFER LIST ENTRY FLAG
+BFLE RBA	EQU	X'80'	ARGUMENT IS RBA
+BFLECHR	EQU	X'40'	ARGUMENT IS CCHHR
+BFLE TTR	EQU	X'20'	ARGUMENT IS TTR
+BFLEAUPD	EQU	X'10'	CVAF UPDATED ARGUMENT FIELD
+BFLEMOD	EQU	X'08'	DATA IN BUF. HAS BEEN MODIFIED
+BFLESKIP	EQU	X'04'	SKIP THIS ENTRY
+BFLEIOER	EQU	X'02'	I/O ERROR
+	DS	XL1	RESERVED
+BFLELTH	DS	XL1	LENGTH OF DSCB BUFFER OR
+*			LENGTH OF VIR DIVIDED BY 256
+BFLEARG	DS	XL5	ARGUMENT OF VIR OR DSCB (CCHHR)
+	ORG	BFLEARG+1	
+BFLEATTR	DS	XL3	'TTR' OF ARGUMENT
+	ORG	BFLEARG+1	
+BFLEARBA	DS	XL4	'RBA' OF ARGUMENT
+BFLEBUF	DS	A	BUFFER ADDRESS
+BFLELN	EQU	*-BFLE	LENGTH OF A BUFFER LIST ENTRY
		IECSDSL1 (1)	FORMAT 1 DSCB DATA SET NAME AND * BUFFER

+IECSDSL1	EQU	*	FORMAT 1 DSCB
+IECSDSF1	EQU	IECSDSL1	FORMAT 1 DSCB
+DS1DSNAM	DS	CL44	DATA SET NAME
+DS1FMTID	DS	CL1	FORMAT IDENTIFIER
+DS1DSSN	DS	CL6	DATA SET SERIAL NUMBER
+DS1VOLSQ	DS	XL2	VOLUME SEQUENCE NUMBER
+DS1CREDT	DS	XL3	CREATION DATE
+DS1EXPDT	DS	XL3	EXPIRATION DATE
+DS1NOEPV	DS	XL1	NUMBER OF EXTENTS ON VOLUME
+DS1NOBDB	DS	XL1	NUMBER OF BYTES USED IN LAST DIRECTORY BLOCK
+*			RESERVED
+	DS	XL1	RESERVED
+DS1SYSCD	DS	CL13	SYSTEM CODE
+	DS	XL7	RESERVED
+DS1DSORG	DS	XL2	DATA SET ORGANIZATION
+*			FIRST BYTE OF DS1DSORG
+DS1DSGIS	EQU	X'80'	IS - INDEXED SEQUENTIAL 001A
+*			ORGANIZATION
+DS1DSGPS	EQU	X'40'	PS - PHYSICAL SEQUENTIAL 001A
+*			ORGANIZATION
+DS1DSGDA	EQU	X'20'	DA - DIRECT ORGANIZATION 001A
+DS1DSGCX	EQU	X'10'	CX - BTAM OR QTAM LINE GROUP 001A
+*			RESERVED 001A
+*			RESERVED 001A

```

+DS1DSGPO EQU X'02' PO - PARTITIONED ORGANIZATION @01A
+DS1DSGU EQU X'01' U - UNMOVABLE, THE DATA @01A
+* CONTAINS LOCATION DEPENDENT
+* INFORMATION
+*
+* SECOND BYTE OF DS1DSORG
+DS1DSGGS EQU X'80' GS - GRAPHICS ORGANIZATION @01A
+DS1DSGTX EQU X'40' TX - TCAM LINE GROUP @01A
+DS1DSGTQ EQU X'20' TQ - TCAM MESSAGE QUEUE @01A
+* EQU X'10' RESERVED @01A
+DS1ACBM EQU X'08' ACCESS METHOD CONTROL BLOCK @01A
+DS1DSGTR EQU X'04' TR - TCAM 3705 @01A
+* EQU X'02' RESERVED @01A
+* EQU X'01' RESERVED @01A
+DS1RECFM DS XL1 RECORD FORMAT
+DS1OPTCD DS XL1 OPTION CODE
+DS1BLKL DS XL2 BLOCK LENGTH
+DS1LRECL DS XL2 RECORD LENGTH
+DS1KEYL DS XL1 KEY LENGTH
+DS1RKP DS XL2 RELATIVE KEY POSITION
+DS1DSIND DS XL1 DATA SET INDICATORS
+DS1SCALO DS XL4 SECONDARY ALLOCATION
+DS1LSTAR DS XL3 LAST USED TRACK AND BLOCK ON TRACK
+DS1TRBAL DS XL2 BYTES REMAINING ON LAST TRACK USED
+ DS XL2 RESERVED
+DS1EXT1 DS XL10 FIRST EXTENT DESCRIPTION
+* FIRST BYTE EXTENT TYPE INDICATOR
+* SECOND BYTE EXTENT SEQUENCE NUMBER
+* THIRD - SIXTH BYTES LOWER LIMIT
+* SEVENTH - TENTH BYTES UPPER LIMIT
+DS1EXT2 DS XL10 SECOND EXTENT DESCRIPTION
+DS1EXT3 DS XL10 THIRD EXTENT DESCRIPTION
+DS1PTRDS DS XL5 POSSIBLE PTR TO A FORMAT 2 OR 3 DSCB
+DS1END EQU *
DSCBLTH EQU *-IECSDSL1-L'DS1DSNAM LENGTH OF DATA PORTION OF DSCB
SAVEAREA DS 18F SAVE AREA
CVPL CVAFSEQ ACCESS=GTEQ, READ DSCB WITH DSN >= SUPPLIED DSN *
IXRCDS=KEEP, KEEP VIERS IN STORAGE DURING CALLS *
DSN=DS1DSNAM, SUPPLIED DATA SET NAME *
BUFLIST=BUFLIST,
MF=L
+ CNOP 0,4
+CVPL EQU *
+ DC CL4'CVPL' EBCDIC 'CVPL'
+ DC AL2(ICV10E-CVPL) LENGTH OF CVPL
+ DC XL1'05' FUNCTION CODE
+ DC XL1'00' STATUS INFORMATION
+ DC B'00001000' FIRST FLAG BYTE
+ DC B'00000000' SECOND FLAG BYTE
+ DC H'0' RESERVED
+ DC A(0) UCB ADDRESS
+ DC A(DS1DSNAM) DATA SET NAME ADDRESS
+ DC A(0) BUFFER LIST ADDRESS
+ DC A(0) INDEX VIR'S BUFFER LIST ADDRESS
+ DC A(0) MAP VIR'S BUFFER LIST ADDRESS
+ DC A(0) I/O AREA ADDRESS
+ DC A(0) DEB ADDRESS
+ DC A(0) ARGUMENT ADDRESS
+ DC A(0) SPACE PARAMETER LIST ADDRESS
+ DC A(0) EXTENT TABLE ADDRESS
+ DC A(0) NEW VRF VIXM BUFFER LIST ADDR
+ DC A(0) VRF DATA ADDRESS
+ DC A(0) COUNT AREA ADDRESS
+ICV10E EQU * END OF CVPL
ORG CVPL EXPAND MAP OVER LIST

```

CVPLMAP ICVAFPL DSECT=NO CVPL MAP

 +* CVAF PARAMETER LIST

+CVPLMAP	DS	OF	CVAF PARAMETER LIST
+CVLBL	DS	CL4	EBCDIC 'CVPL'
+CVLTH	DS	H	LENGTH OF CVPL
+CVFCTN	DS	XL1	FUNCTION BYTE
+CVDIRD	EQU	X'01'	CVAFDIR ACCESS=READ
+CVDIWR	EQU	X'02'	CVAFDIR ACCESS=WRITE
+CVDIRLS	EQU	X'03'	CVAFDIR ACCESS=RLSE
+CVSEQGT	EQU	X'04'	CVAFSEQ ACCESS=GT
+CVSEQGTE	EQU	X'05'	CVAFSEQ ACCESS=GTEQ
+CVDMIXA	EQU	X'06'	CVAFDSM ACCESS=IXADD
+CVDMIXD	EQU	X'07'	CVAFDSM ACCESS=IXDLT
+CVDMALC	EQU	X'08'	CVAFDSM ACCESS=ALLOC
+CVDMRLS	EQU	X'09'	CVAFDSM ACCESS=RLSE
+CVDMMAP	EQU	X'0A'	CVAFDSM ACCESS=MAPDATA
+CVVOL	EQU	X'0B'	CVAFVOL ACCESS=VIBBLD
+CVVRFRD	EQU	X'0C'	CVAFVRF ACCESS=READ
+CVVRFWR	EQU	X'0D'	CVAFVRF ACCESS=WRITE
+CVSTAT	DS	XL1	STATUS INFORMATION (SEE LIST * BELOW)
+CVFL1	DS	XL1	FIRST FLAG BYTE
+CV1IIVT	EQU	X'80'	INDEXED VTOC ACCESSED
+CV1IOAR	EQU	X'40'	IOAREA=KEEP
+CV1PGM	EQU	X'20'	BRANCH=(YES,PGM)
+CV1MRCDS	EQU	X'10'	MAPRCDS=YES
+CV1IRCDS	EQU	X'08'	IXRCDS=KEEP
+CV1MAPIX	EQU	X'04'	MAP=INDEX
+CV1MAPVT	EQU	X'02'	MAP=VTOC
+CV1MAPVL	EQU	X'01'	MAP=VOLUME
+CVFL2	DS	XL1	SECOND FLAG BYTE
+CV2HIVIE	EQU	X'80'	HIVIER=YES
+CV2VRF	EQU	X'40'	VRF DATA EXISTS
+CV2CNT	EQU	X'20'	COUNT=YES
+CV2RCVR	EQU	X'10'	RECOVER=YES
+CV2SRCH	EQU	X'08'	SEARCH=YES
+CV2DSNLY	EQU	X'04'	DSNONLY=YES
+CV2VER	EQU	X'02'	VERIFY=YES
+CV2NLEVL	EQU	X'01'	OUTPUT-NEW HIGHEST LEVEL VIER CREATED
+*			RESERVED
+CVUCB	DS	A	UCB ADDRESS
+CVDSN	DS	A	DATA SET NAME ADDRESS
+CVBUFL	DS	A	BUFFER LIST ADDRESS
+CVIRCDS	DS	A	INDEX VIR'S BUFFER LIST ADDRESS
+CVMRCDS	DS	A	MAP VIR'S BUFFER LIST ADDRESS
+CVIOAR	DS	A	I/O AREA ADDRESS
+CVDEB	DS	A	DEB ADDRESS
+CVARG	DS	A	ARGUMENT ADDRESS
+CVSPACE	DS	A	SPACE PARAMETER LIST ADDRESS
+CVEXTS	DS	A	EXTENT TABLE ADDRESS
+CVBUFL2	DS	A	NEW VRF VIXM BUFFER LIST ADDR
+CVVRFDA	DS	A	VRF DATA ADDRESS
+CVCTAR	DS	A	COUNT AREA ADDRESS
+CVPLNGTH	EQU	*-CVPLMAP	

+* VALUES OF CVSTAT
 +*(THIS PART OF THE ICVAFPL MACRO EXPANSION IS NOT SHOWN)
 END

EXAMPLE 4: USING THE CVAFSEQ MACRO WITH A NONINDEXED VTOC

This example reads as many as five DSCBs in physical-sequential order. The address of the UCB is supplied to the program in register 5 (labeled RUCB). The address of a parameter list is supplied in register 4 (labeled RLIST). The first word of the parameter list contains the address of a 5-byte field. On entry, this field is set to zero if no previous DSCBs have been read; otherwise, the field is set to the CCHHR of the last DSCB read. This 5-byte field is supplied by the caller of this program and is not modified by this program.

The remainder of the parameter list consists of one or more 2-word entries, to a maximum of five 2-word entries. The first word of each entry contains the address of a 140-byte DSCB buffer. The second word contains the address of a 5-byte field that is to contain the CCHHR of the DSCB.

A buffer list with five buffer list entries is contained in the program. The ICVAFBFL macro generates the buffer list header and one buffer list entry. The remaining buffer list entries are generated following the ICVAFBFL macro.

The CVAFSEQ macro is used once in the program to read as many DSCBs as there are 2-word entries in the parameter list. The buffer list header field BFLHNOE is initialized with the number of buffer list entries that CVAFSEQ is to process. The number matches the number of 2-word entries in the parameter list supplied to this program.

After the CVAFSEQ call, the CCHHR for each DSCB read is moved from the buffer list entry field BFLEARG to the field whose address is supplied by the caller of the program. If the BFLEARG field is zero, the previous DSCB read was the last in the VTOC.

The BFLEARG in the first buffer list entry is initialized with the CCHHR supplied by the caller: its address is the third word in the parameter list. This CCHHR serves as the starting place for the CVAFSEQ call. DSCBs with a CCHHR greater than the supplied CCHHR are read.

This program must be APF authorized.

```

EXAMPLE4 CSECT
STM      14,12,12(13)
BALR    12,0
USING   *,12
ST      13,SAVEAREA+4
LA      RWORK,SAVEAREA
ST      RWORK,8(,13)
LR      13,RWORK
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*      REGISTERS
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
REG1     EQU    1      REGISTER 1
RWORK    EQU    3      WORK REGISTER
RLIST    EQU    4      ADDRESS OF PARM LIST
RUCB     EQU    5      UCB ADDRESS
RCURRENT EQU    6      CURRENT ENTRY IN PARM LIST
RBLE     EQU    7      CURRENT BUFFER LIST ENTRY
RCOUNT   EQU    8      COUNT OF ENTRIES IN BUFFER LIST
REG15    EQU    15     RETURN CODE REGISTER 15

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
* READ UP TO 5 DSCBS.
* RUCB CONTAINS ADDRESS OF UCB.
* RLST CONTAINS ADDRESS OF PARAMETER LIST.
* WORD 0 = ADDRESS OF CCHHR OF LAST DSCB READ. THIS DSCB IS
* NOT TO BE READ
* WORD 1 = ADDRESS OF DSCB BUFFER.
* WORD 2 = ADDRESS OF CCHHR OF DSCB READ.
* WORD1 AND WORD2 REPEATED UP TO 4 TIMES.
* HIGH ORDER BIT OF WORD 2 SET TO ONE FOR LAST ENTRY.
*

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
* USING LIST,RLIST ADDRESSABILITY TO PARM LIST
XC BFLHDR(BFLHLN+5*BFLLELN),BFLHDR ZERO BUFFER LIST WITH *
* 5 BUFFER LIST ENTRIES
OI BFLHFL,BFLHDSCB DSCBS TO BE READ WITH BUFFER LIST
LA RCURRENT,LISTPRMS FIRST DOUBLEWORD ENTRY IN PARM LIST
USING LISTPRMS,RCURRENT USING ON DOUBLEWORDS
LA RBLE,BFLE FIRST BUFFER LIST ENTRY
USING BFLE,RBLE
L RWORK,LISTSTRT ADDRESS OF STARTING CCHHR
MVC BFLEARG,0(RWORK) MOVE STARTING CCHHR INTO FIRST *
* BUFFER LIST ENTRY
XR RCOUNT,RCOUNT ZERO COUNT
BUFLOOP EQU * PUT BUFFER ADDRESSES IN BUFFER LIST *
* ENTRIES
LA RCOUNT,1(RCOUNT) INCREMENT COUNT
L RWORK,LISTBUF ADDRESS OF DSCB BUFFER
ST RWORK,BFLEBUF-BFLE(,RBLE) PLACE IN BUFFER LIST
MVI BFLELTH-BFLE(RBLE),DSCBLTH FULL DSCB READ
TM LISTLAST,LASTBIT IS IT LAST ENTRY IN LIST
LA RCURRENT,LISTNEXT INCREMENT TO NEXT ENTRY IN LIST
LA RBLE,BFLELN(,RBLE) INCREMENT TO NEXT BUFFER LIST ENTRY
BZ BUFLOOP LOOP TO PUT NEXT BUFFER IN BFLE
STC RCOUNT,BFLHNOE SET NUMBER OF ENTRIES IN BUFFER *
* LIST HEADER
DROP RCURRENT,RBLE

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
* READ UP TO 5 DSCBS WHOSE CCHHR IS GREATER THAN THE CCHHR IN
* THE FIRST BUFFER LIST ENTRY
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
* CVAFSEQ UCB=(RUCB), CALL CVAF *
* BRANCH=YES, BRANCH ENTER *
* MF=(E,CVPL)
+ LA 1,CVPL LOAD PARAMETER REG 1
+ ST RUCB,12(,1) STORE UCB PTR IN PARM LIST
+ L 15,16 LOAD THE CVT
+ L 15,328(,15) LOAD VS1/VS2 COMMON EXTENSION2
+ L 15,12(,15) LOAD THE CVAF TABLE ADDRESS
+ L 15,0(,15) LOAD THE CVAF ADDRESS
+ BALR 14,15 BRANCH AND LINK TO CVAF
USING CVPL,REG1 ADDRESSABILITY TO CVPL
LTR REG15,REG15 ANY ERROR
BZ MOVECHR BRANCH IF MOVE IN CCHHRS
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

```

*
* DETERMINE WHAT ERROR IS
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
* C REG15,ERROR4 IS RETURN CODE 4
* BNE OTHERERR BRANCH IF NOT 4
* CLI CVSTAT,STAT032 IS IT END OF DATA?
* BNE OTHERERR BRANCH IF NOT
* DROP REG1 ADDRESSABILITY TO CVPL NOT NEEDED

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*      DETERMINE IF ANY DSCBS HAVE BEEN READ. BFLEARG IS NON-ZERO
*      IN EACH BUFFER LIST ENTRY FOR WHICH A DSCB HAS BEEN READ
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
MOVECHR EQU *      IS DATA SET NAME GREATER THAN DSN2
        LA RCURRENT,LISTPRMS      FIRST ENTRY IN PARM LIST
        USING LISTPRMS,RCURRENT
        LA RBLE,BFLE      FIRST BUFFER LIST ENTRY
        USING BFLE,RBLE
CHRLOOP EQU *      MOVE CCHHR ARGUMENT TO CALLER AREA
        L RWORK,LISTCHR      ADDRESS OF CCHHR OF CALLER
        XC 0(L'BFLEARG,RWORK),0(RWORK) ZERO CALLER CCHHR AREA
        NC BFLEARG,BFLEARG      IS CCHHR ZERO
        BZ EXIT      BRANCH IF YES-NO MORE DSCBS
        MVC 0(L'BFLEARG,RWORK),BFLEARG MOVE CCHHR TO CALLER AREA
        TM LISTLAST,LASTBIT      LAST ENTRY IN PARM LIST?
        BO EXIT      BRANCH IF YES
        LA RCURRENT,LISTNEXT      NEXT ENTRY IN LIST
        LA RBLE,BFLELN(,RBLE)      NEXT BUFFER LIST ENTRY
        B CHRLOOP      TEST NEXT BFLE
EXIT EQU *      RETURN TO CALLER
        L 13,SAVEAREA+4
        RETURN (14,12)
+      LM 14,12,12(13)      RESTORE THE REGISTERS
+      BR 14      RETURN

OTHERERR EQU *      ERROR PROCESSING
*
*
*
        B EXIT      RETURN
ERROR4 DC F'4'      RETURN CODE 4
        ICVAFBFL DSECT=NO      BUFFER LIST WITH ONE BUFFER LIST *
        ENTRY

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
+*      BUFFER LIST HEADER
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

```

+BFLHDR DS OF      BUFFER LIST HEADER
+BFLHNOE DS XL1      NUMBER OF ENTRIES
+BFLHFL DS XL1      KEY AND FLAG BYTE
+      ORG BFLHFL
+BFLHKEY DS XL1      PROTECT KEY (FIRST 4 BITS)
+BFLHVIR EQU X'08'      BUF. LIST ENTRIES DESCRIBE VIRS
+BFLHDSCB EQU X'04'      BUF. LIST ENTRIES DESCRIBE DSCBS
+      DS XL1      RESERVED
+BFLHSP DS XL1      SUBPOOL OF BUF. LIST/BUFFERS
+BFLHFCHN DS A      FORWARD CHAIN PTR TO NEXT BUF.
+*      LIST
+BFLHLN EQU *-BFLHDR      LENGTH OF BUFFER LIST HEADER

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
+*      BUFFER LIST ENTRY
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

```

+BFLE DS OF      BUFFER LIST ENTRY
+BFLEFL DS XL1      BUFFER LIST ENTRY FLAG
+BFLELBA EQU X'80'      ARGUMENT IS RBA
+BFLECHR EQU X'40'      ARGUMENT IS CCHHR
+BFLELTR EQU X'20'      ARGUMENT IS TTR
+BFLEAUPD EQU X'10'      CVAF UPDATED ARGUMENT FIELD
+BFLEMOD EQU X'08'      DATA IN BUF. HAS BEEN MODIFIED
+BFLESKIP EQU X'04'      SKIP THIS ENTRY
+BFLEIOER EQU X'02'      I/O ERROR
+      DS XL1      RESERVED
+BFLELTH DS XL1      LENGTH OF DSCB BUFFER OR

```



```

+DS1PTRDS DS XL5 POSSIBLE PTR TO A FORMAT 2 OR 3 DSCB
+DS1END EQU *
DSCBLTH EQU *-IECSDSL1 LENGTH OF DSCB
LIST DSECT * PARAMETER LIST
LISTSTRT DS F ADDRESS OF CCHHR TO START SEARCH
LISTPRMS EQU *
LISTBUF DS F BUFFER ADDRESS
LISTCHR DS 0F ADDRESS OF CCHHR FIELD
LISTLAST DS X BYTE
LASTBIT EQU X'80' LAST DOUBLE WORD
DS AL3 3 BYTE ADDRESS OF CCHHR
LISTNEXT EQU * NEXT DOUBLEWORD
EXAMPLE4 CSECT

```

```

*
* READ DSCBS WITH CCHHR GREATER THAN THE CCHHR IN THE FIRST
* BUFFER LIST ENTRY.
*
*****

```

```

CVPL CVAFSEQ ACCESS=GT, ADDRESS OF BUFFER LIST *
      BUFLIST=BFLHDR, *
      MF=L
+ CNOP 0,4
+CVPL EQU *
+ DC CL4'CVPL' EBCDIC 'CVPL'
+ DC AL2(ICV6E-CVPL) LENGTH OF CVPL
+ DC XL1'04' FUNCTION CODE
+ DC XL1'00' STATUS INFORMATION
+ DC B'00100000' FIRST FLAG BYTE
+ DC B'00000000' SECOND FLAG BYTE
+ DC H'0' RESERVED
+ DC A(0) UCB ADDRESS
+ DC A(0) DATA SET NAME ADDRESS
+ DC A(BFLHDR) BUFFER LIST ADDRESS
+ DC A(0) INDEX VIR'S BUFFER LIST ADDRESS
+ DC A(0) MAP VIR'S BUFFER LIST ADDRESS
+ DC A(0) I/O AREA ADDRESS
+ DC A(0) DEB ADDRESS
+ DC A(0) ARGUMENT ADDRESS
+ DC A(0) SPACE PARAMETER LIST ADDRESS
+ DC A(0) EXTENT TABLE ADDRESS
+ DC A(0) NEW VRF VIXM BUFFER LIST ADDR
+ DC A(0) VRF DATA ADDRESS
+ ICV6E EQU * END OF CVPL
      ORG CVPL EXPAND MAP OVER LIST
CVPLMAP ICVAFPL DSECT=NO CVPL MAP

```

+* CVAF PARAMETER LIST

```

+CVPLMAP DS 0F CVAF PARAMETER LIST
+CVLBL DS CL4 EBCDIC 'CVPL'
+CVLTH DS H LENGTH OF CVPL
+CVFCTN DS XL1 FUNCTION BYTE
+CVDIRD EQU X'01' CVAFDIR ACCESS=READ
+CVDIWR EQU X'02' CVAFDIR ACCESS=WRITE
+CVDIRLS EQU X'03' CVAFDIR ACCESS=RLSE
+CVSEQGT EQU X'04' CVAFSEQ ACCESS=GT
+CVSEQGTE EQU X'05' CVAFSEQ ACCESS=GTEQ
+CVDMIXA EQU X'06' CVAFDSM ACCESS=IXADD
+CVDMIXD EQU X'07' CVAFDSM ACCESS=IXDLT
+CVDMALC EQU X'08' CVAFDSM ACCESS=ALLOC
+CVDMRLS EQU X'09' CVAFDSM ACCESS=RLSE
+CVDMMAP EQU X'0A' CVAFDSM ACCESS=MAPDATA
+CVVOL EQU X'0B' CVAFVOL ACCESS=VIBBLD
+CVVFRD EQU X'0C' CVAFVRF ACCESS=READ

```

+CVVRFWR	EQU	X'0D'	CVAFVRF ACCESS=WRITE
+CVSTAT	DS	XL1	STATUS INFORMATION (SEE LIST * BELOW)
+			FIRST FLAG BYTE
+CVFL1	DS	XL1	INDEXED VTOC ACCESSED
+CV1IIVT	EQU	X'80'	IOAREA=KEEP
+CV1IOAR	EQU	X'40'	BRANCH=(YES,PGM)
+CV1PGM	EQU	X'20'	MAPRCD=YES
+CV1MRCDS	EQU	X'10'	IXRCD=KEEP
+CV1IRCDS	EQU	X'08'	MAP=INDEX
+CV1MAPIX	EQU	X'04'	MAP=VTOC
+CV1MAPVT	EQU	X'02'	MAP=VOLUME
+CV1MAPVL	EQU	X'01'	SECOND FLAG BYTE
+CVFL2	DS	XL1	HIVIER=YES
+CV2HIVIE	EQU	X'80'	VRF DATA EXISTS
+CV2VRF	EQU	X'40'	COUNT=YES
+CV2CNT	EQU	X'20'	RECOVER=YES
+CV2RCVR	EQU	X'10'	SEARCH=YES
+CV2SRCH	EQU	X'08'	DSNONLY=YES
+CV2DSNLY	EQU	X'04'	VERIFY=YES
+CV2VER	EQU	X'02'	OUTPUT-NEW HIGHEST LEVEL VIER
+CV2NLEVL	EQU	X'01'	CREATED
+*			RESERVED
+	DS	H	UCB ADDRESS
+CVUCB	DS	A	DATA SET NAME ADDRESS
+CVDSN	DS	A	BUFFER LIST ADDRESS
+CVBUFL	DS	A	INDEX VIR'S BUFFER LIST ADDRESS
+CVIRCDS	DS	A	MAP VIR'S BUFFER LIST ADDRESS
+CVMRCDS	DS	A	I/O AREA ADDRESS
+CVIOAR	DS	A	DEB ADDRESS
+CVDEB	DS	A	ARGUMENT ADDRESS
+CVARG	DS	A	SPACE PARAMETER LIST ADDRESS
+CVSPACE	DS	A	EXTENT TABLE ADDRESS
+CVEXTS	DS	A	NEW VRF VIXM BUFFER LIST ADDR
+CVBUFL2	DS	A	VRF DATA ADDRESS
+CVVRFDA	DS	A	COUNT AREA ADDRESS
+CVCTAR	DS	A	
+CVPLNGTH	EQU	*-CVPLMAP	

+* VALUES OF CVSTAT
+*(THIS PART OF THE ICVAFPL MACRO EXPANSION IS NOT SHOWN)
END

EXAMPLE 5: USING THE CVAFTST AND CVAFDSM MACROS

This example returns a format-5 DSCB to the caller. The format-5 DSCB is constructed by this program if the volume contains an indexed VTOC. The format-5 DSCB is read by another program, F5RTN (not described in the example), if the volume contains a nonindexed VTOC.

The CVAFTST macro is used to determine if a nonindexed VTOC is on the volume.

If the CVAFTST return code is neither 0 nor 4 (a nonindexed VTOC is on the volume), the CVAFDSM macro is issued to obtain up to 27 extents from the VPSM in the VTOC index. The program does not determine whether the CVAFTST return code is 8 (volume contains indexed VTOC) or 12 (it cannot be determined what type of VTOC is on the volume). In either case, the CVAFDSM macro is issued. If the CVAFTST return code is 12, the CVAFDSM macro call will cause CVAF to determine whether an indexed or a nonindexed VTOC is on the volume, and the CV1IVT bit will be set to one or zero, accordingly.

The extent table (at label EXTABL) is initialized to request 27 extents from the CVAFDSM macro, which is one more than the number of extents that fit in a format-5 DSCB. The format-5 DSCB is constructed from the first 26 extents returned from the CVAFDSM call.

The first extent in the extent table is initialized from the last extent in the format-5 DSCB area supplied by the caller of the program. If this is the first call, the program assumes that the format-5 area is initialized to zero. Thus, the first extent in the extent table has a value of zero to serve as the starting place for the extent search. If this is the second or subsequent call, the last extent in the format-5 area would be the last extent obtained from the previous CVAFDSM call.

The format-5 chain pointer field (DS5PTRDS) is set to a nonzero value if CVAFDSM returned a 27th extent. In this case, the program will be called again to obtain another format-5 DSCB.

The program's return code is 0 if no errors were encountered and 4 if an error was encountered.

This program must be APF authorized.

```
EXAMPLE5 CSECT
          STM 14,12,12(13)
          BALR 12,0
          USING *,12
          ST 13,SAVEAREA+4
          LA RWORK,SAVEAREA
          ST RWORK,8(,13)
          LR 13,RWORK
```

XX

⌘

REGISTERS

⌘

XX

```
RDEB EQU 3 DEB ADDRESS SUPPLIED BY CALLER
RUCB EQU 4 UCB ADDRESS SUPPLIED BY CALLER
RF5 EQU 5 ADDRESS OF FORMAT 5 BUFFER SUPPLIED BY CALLER
RWORK EQU 6 WORK REGISTER
REG15 EQU 15 RETURN CODE REGISTER 15
⌘
KF5 EQU 26 NUMBER OF FORMAT 5 EXTENTS
```

*
* READ FORMAT 5 DSCB OR BUILD A FORMAT 5 DSCB IF
* AN INDEXED VTOC
* UCB ADDRESS SUPPLIED IN RUCB.
* RF5 CONTAINS THE ADDRESS OF THE FORMAT 5 DSCB BUFFER. IT
* CONTAINS THE LAST FORMAT 5 DSCB READ OR BUILT. THE FORMAT 5
* BUFFER IS ZERO IF THIS IS THE FIRST CALL
* IF THE FORMAT 5 DSCB BUFFER RETURNED TO THE CALLER HAS A
* NONZERO VALUE IN DS5PTRDS, THIS ROUTINE WILL BE CALLED
* AGAIN TO OBTAIN THE NEXT FORMAT 5 DSCB.
*

	USING IECSDSF5,RF5	ADDRESSABILITY TO FORMAT 5 BUFFER
	CVAFSTST UCB=(RUCB)	TEST VTOC
	CNOP 0,4	START OF CVAFSTST MACRO
	LR 1,RUCB	LOAD PARAMETER REG 1
+	L 15,16	LOAD THE CVT
+	L 15,328(,15)	LOAD VS1/V52 COMMON EXTENSION2
+	L 15,12(,15)	LOAD THE CVAF TABLE ADDRESS
+	LTR 15,15	TEST FOR ZERO VALUE
+	BZ ICVIE	CVAF IS NOT ON THE SYSTEM
+	L 15,4(,15)	LOAD POINTER TO CVAF TEST E.P.
+	BALR 14,15	BRANCH AND LINK TO CVAF TEST
+ICVIE	EQU *	END OF CVAFSTST
	LTR REG15,REG15	
	BZ UNINDEXD	READ NEXT FORMAT 5
	C REG15,NOTIXRC	UNINDEXED VTOC?
	BE UNINDEXD	READ NEXT FORMAT 5

*
* ASSUME INDEXED VTOC UNLESS CVAFDSM CALL INDICATES UNINDEXED
*

	MVC EXTS(L'DS5AVEXT),DS5MAVET+L'DS5MAVET-L'DS5AVEXT	MOVE THE	*
		LAST EXTENT FROM FORMAT 5 TO FIRST	*
		ENTRY IN THE EXTENT TABLE	
	CVAFDSM MF=(E,CVPL),	GET 27 EXTENTS FROM CVPL	*
	UCB=(RUCB),	RUCB ADDRESS REQUIRED	*
	DEB=(RDEB),	RDEB ADDRESS REQUIRED BY	*
		UNAUTHORIZED PROGRAMS CALLING CVAF	*
	BRANCH=YES	BRANCH ENTRY CALL	*
+	LA 1,CVPL	LOAD PARAMETER REG 1	
+	L 15,16	LOAD THE CVT	
+	L 15,328(,15)	LOAD VS1/V52 COMMON EXTENSION2	
+	L 15,12(,15)	LOAD THE CVAF TABLE ADDRESS	
+	L 15,0(,15)	LOAD THE CVAF ADDRESS	
+	BALR 14,15	BRANCH AND LINK TO CVAF	
	TM CVFL1,CV1IVT	IS THIS INDEXED VTOC	
	BZ UNINDEXD	READ FORMAT 5 IF NOT	
	LTR REG15,REG15	ANY ERROR	
	BZ NOERROR		
	C REG15,RC04		
	BNE OTHERERR	UNEXPECTED ERROR	
	CLI CVSTAT,STAT032	END OF DATA	
	BNE OTHERERR	UNEXPECTED ERROR	
NOERROR	EQU *	BUILD FORMAT 5	
	MVC DS5KEYID,F5ID		
	MVC DS5AVEXT(L'DS5AVEXT+L'DS5EXTAV),EXTS	MOVE IN EXTENTS	*
		TO DS5FMTID	
	MVI DS5FMTID,C'5'		
	MVC DS5MAVET,EXTS+L'DS5AVEXT+L'DS5EXTAV	MOVE REMAINING	*
		EXTENTS	
	XR REG15,REG15	RETURN CODE ZERO	
	XC DS5PTRDS,DS5PTRDS	ZERO CHAIN POINTER	
	NC EXTS+L'EXTS-L'DS5AVEXT(L'DS5AVEXT),EXTS+L'EXTS-L'DS5AVEXT	IS LAST(27TH) EXTENT FROM CVAF	*
		ZERO?	

```

BZ RETURN BRANCH IF YES-LEAVE DS5PTRDS ZERO
MVI DS5PTRDS+L'DS5PTRDS-1,1 SET DS5PTRDS NONZERO TO SIMULATE *
THERE BEING ANOTHER FORMAT 5

UNINDXD B RETURN CALL ROUTINE TO READ NEXT FORMAT 5
EQU * LINK EP=F5RTN LINK TO FORMAT 5 ROUTINE. RETURN *
CODE PASSED BACK IN REG15

+ CNOP 0,4
+ BAL 15,*+20 LOAD SUP.PARAMLIST ADR
+ DC A(*+8) ADDR OF EP PARAMETER
+ DC A(0) DCB ADDRESS PARAMETER LCOA
+ DC CL8'F5RTN' EP PARAMETER
+ SVC 6 ISSUE LINK SVC
RETURN EQU * RETURN TO CALLER
L 13,SAVEAREA+4
RETURN (14,12),RC=(15)
+ L 14,12(13,0) RESTORE REGISTER 14
+ LM 0,12,20(13) RESTORE THE REGISTERS
+ BR 14 RETURN
OTHERERR EQU * ERROR
L REG15,RC04 ERROR RETURN CODE
B RETURN

DSCB DSECT
IECSDSL1 (5)
+IECSDSL5 EQU * FORMAT 5 DSCB
+IECSDSF5 EQU IECSDSL5
+DS5KEYID DS XL4 KEY IDENTIFIER
+DS5AVEXT DS XL5 AVAILABLE EXTENT
+* BYTES 1 - 2 RELATIVE TRACK ADDRESS OF THE FIRST TRACK
+* IN THE EXTENT
+* BYTES 3 - 4 NUMBER OF UNUSED CYLINDERS IN THE EXTENT
+* BYTE 5 NUMBER OF ADDITIONAL UNUSED TRACKS
+DS5EXTAV DS XL35 SEVEN AVAILABLE EXTENTS
+DS5FMTID DS CL1 FORMAT IDENTIFIER
+DS5MAVET DS XL90 EIGHTEEN AVAILABLE EXTENTS
+DS5PTRDS DS XL5 POINTER TO NEXT FORMAT 5 DSCB
+DS5END EQU *
EXAMPLE5 CSECT
NOTIXRC DC F'4' CVAFTST RETURN CODE-UNINDEXED
RC04 DC F'4' RETURN CODE 4
F5ID DC XL4'0505050505' FORMAT 5 FIELD, DS5KEYID
SAVEAREA DS 18F REGISTER SAVE AREA
EXTABL DS 0CL(1+(KF5+1)*L'DS5AVEXT) EXTENT TABLE
EXTNO DC AL1(KF5+1) NUMBER OF EXTENTS IN TABLE
EXTS DS CL((KF5+1)*L'DS5AVEXT) EXTENTS
CVPL CVAFDSM ACCESS=MAPDATA, *
COUNT=NO, DO NOT COUNT EXTENTS *
MAP=VOLUME, ACCESS VOLUME SPACE MAP *
EXTENTS=EXTABL, EXTENT TABLE ADDRESS *
MF=L LIST FORM OF MACRO *

+ CNOP 0,4
+CVPL EQU *
+ DC CL4'CVPL' EBCDIC 'CVPL'
+ DC AL2(ICV9E-CVPL) LENGTH OF CVPL
+ DC XL1'0A' FUNCTION CODE
+ DC XL1'00' STATUS INFORMATION
+ DC B'001000001' FIRST FLAG BYTE
+ DC B'000000000' SECOND FLAG BYTE
+ DC H'0' RESERVED
+ DC A(0) UCB ADDRESS
+ DC A(0) DATA SET NAME ADDRESS
+ DC A(0) BUFFER LIST ADDRESS
+ DC A(0) INDEX VIR'S BUFFER LIST ADDRESS
+ DC A(0) MAP VIR'S BUFFER LIST ADDRESS
+ DC A(0) I/O AREA ADDRESS
+ DC A(0) DEB ADDRESS

```

```

+          DC      A(0)          ARGUMENT ADDRESS
+          DC      A(0)          SPACE PARAMETER LIST ADDRESS
+          DC      A(EXTABL)     EXTENTS TABLE ADDRESS
+          DC      A(0)          NEW VRF VIXM BUFFER LIST ADDR
+          DC      A(0)          VRF DATA ADDRESS
+          DC      A(0)          COUNT AREA ADDRESS
+ICV9E     EQU      *           END OF CVPL
          ORG      CVPL         OVERLAY CVPL WITH EXPANSION OF MAP
          CVPLMAP ICVAFPL DSECT=NO
+*****
+*          CVAF PARAMETER LIST
+*****
+CVPLMAP   DS      OF           CVAF PARAMETER LIST
+CVLBL     DS      CL4         EBCDIC 'CVPL'
+CVLTH     DS      H           LENGTH OF CVPL
+CVFCTN    DS      XL1        FUNCTION BYTE
+CVDIRD    EQU     X'01'       CVAFDIR ACCESS=READ
+CVDIWR    EQU     X'02'       CVAFDIR ACCESS=WRITE
+CVDIRLS   EQU     X'03'       CVAFDIR ACCESS=RLSE
+CVSEQGT   EQU     X'04'       CVAFSEQ ACCESS=GT
+CVSEQGTE  EQU     X'05'       CVAFSEQ ACCESS=GTEQ
+CVDMIXA   EQU     X'06'       CVAFDSM ACCESS=IXADD
+CVDMIXD   EQU     X'07'       CVAFDSM ACCESS=IXDLT
+CVDMALC   EQU     X'08'       CVAFDSM ACCESS=ALLOC
+CVDMRLS   EQU     X'09'       CVAFDSM ACCESS=RLSE
+CVDMMAP   EQU     X'0A'       CVAFDSM ACCESS=MAPDATA
+CVVOL     EQU     X'0B'       CVAFVOL ACCESS=VIBBLD
+CVVRFRD   EQU     X'0C'       CVAFVRF ACCESS=READ
+CVVRFWR   EQU     X'0D'       CVAFVRF ACCESS=WRITE
+CVSTAT    DS      XL1        STATUS INFORMATION (SEE LIST      X
+          +           BELOW)
+CVFL1     DS      XL1        FIRST FLAG BYTE
+CV1IVT    EQU     X'80'       INDEXED VTOC ACCESSED
+CV1IOAR   EQU     X'40'       IOAREA=KEEP
+CV1PGM    EQU     X'20'       BRANCH=(YES,PGM)
+CV1MRCDS  EQU     X'10'       MAPRCDS=YES
+CV1IRCDS  EQU     X'08'       IXRCDS=KEEP
+CV1MAPIX  EQU     X'04'       MAP=INDEX
+CV1MAPVT  EQU     X'02'       MAP=VTOC
+CV1MAPVL  EQU     X'01'       MAP=VOLUME
+CVFL2     DS      XL1        SECOND FLAG BYTE
+CV2HIVIE  EQU     X'80'       HIVIER=YES
+CV2VRF    EQU     X'40'       VRF DATA EXISTS
+CV2CNT    EQU     X'20'       COUNT=YES
+CV2RCVR   EQU     X'10'       RECOVER=YES
+CV2SRCH   EQU     X'08'       SEARCH=YES
+CV2DSNLY  EQU     X'04'       DSNONLY=YES
+CV2VER    EQU     X'02'       VERIFY=YES
+CV2NLEVL  EQU     X'01'       OUTPUT-NEW HIGHEST LEVEL VIER
+*          +           CREATED
+          +           RESERVED
+CVUCB     DS      A           UCB ADDRESS
+CVDSN     DS      A           DATA SET NAME ADDRESS
+CVBUFL    DS      A           BUFFER LIST ADDRESS
+CVIRCDS   DS      A           INDEX VIR'S BUFFER LIST ADDRESS
+CVMRCDS   DS      A           MAP VIR'S BUFFER LIST ADDRESS
+CVIOAR    DS      A           I/O AREA ADDRESS
+CVDEB     DS      A           DEB ADDRESS
+CVARG     DS      A           ARGUMENT ADDRESS
+CVSPACE   DS      A           SPACE PARAMETER LIST ADDRESS
+CVEXTS    DS      A           EXTENT TABLE ADDRESS
+CVBUFL2   DS      A           NEW VRF VIXM BUFFER LIST ADDR
+CVVRFDA   DS      A           VRF DATA ADDRESS
+CVCTAR    DS      A           COUNT AREA ADDRESS
+CVPLNGTH  EQU     *-CVPLMAP
+*          +           VALUES OF CVSTAT
+*(THIS PART OF THE ICVAFPL MACRO EXPANSION IS NOT SHOWN)
          END

```

APPENDIX C. VTOC INDEX ERROR MESSAGE AND ASSOCIATED CODES

ERROR MESSAGE

When CVAF finds an error in a VTOC index, it issues this message:

```
IEC606I VTOC INDEX DISABLED ON dev,volser,  
code,[rba[,secno,offset]]
```

In addition, CVAF puts a return code in the CVSTAT field of the CVPL.

EXPLANATION

The Common VTOC Access Facility (CVAF) detected a VTOC index error on the device 'dev' with volume serial number 'volser'. 'code' is a number that represents the kind of VTOC index error encountered. 'rba' is the RBA of the VIR in the VTOC index that contains a structure error indicated by 'code'. If the VIR is a VIER, the section number in the VIER containing the VTOC index entry is supplied in 'secno', and the offset into the section of that VTOC index entry is supplied in 'offset'.

SYSTEM ACTION

The VTOC index is disabled by zeroing the index bit in the format-4 DSCB and setting the bit in the first high-level VIER that indicates invalid VTOC index structure. The VTOC will be converted to nonindexed format when DADSM next allocates space on the volume. A system dump is written to the SYS1.DUMP data set, and an entry is made in the SYS1.LOGREC data set. The message IEC604I (which indicates that the VTOC convert routines have been used) will be issued later.

PROGRAMMER RESPONSE

Examine the system dump and a print of the VTOC index, and use the information in message IEC606I to determine the cause of the VTOC index structure error.

ROUTING AND DESCRIPTOR CODES

The routing codes are 4 (direct access pool) and 10 (system/error maintenance), and the descriptor code is 4 (system status).

CODES PUT IN THE CVSTAT FIELD

Code	Meaning	Code	Meaning
0(00)	No error.	25(19)	Conflicting or incomplete information specified in the space table for a CVAFDSM ACCESS=ALLOC, MAP=VOLUME request.
1(01)	Data set name not found, or VIER is empty.	27(1B)	VTOC index full. No free VIRs available and a VIER split is required.
2(02)	Argument is outside VTOC extents or RBA range of VTOC index.	28(1C)	Space keyword omitted (CVSPACE field zero in CVPL).
4(04)	Invalid parameter supplied (wrong key).	29(1D)	CVAFDSM ACCESS=ALLOC: No format 0 DSCB available (MAP=VTOC), or VTOC index full (MAP=INDEX), or volume space not available (MAP=VOLUME).
5(05)	DSN keyword omitted.	30(1E)	CVAFDSM ACCESS=ALLOC: CCHHR (MAP=VTOC) or RBA MAP=INDEX or volume space extent (MAP=VOLUME) already allocated.
6(06)	Not authorized to perform this function.	31(1F)	CVAFDSM ACCESS=ALLOC or ACCESS=MAPDATA: CCHHR supplied outside VTOC extents (MAP=VTOC), or RBA outside VTOC index extents (MAP=INDEX), or volume space extent invalid or outside volume (MAP=VOLUME).
7(07)	Buffer list omitted.	32(20)	End of data. CVAFDSM ACCESS=MAPDATA: no more free extents in VPSM. CVAFSEQ: no more names in index or DSCBs in VTOC. For indexed access, no DSN in VTOC index with higher or higher-or-equal key than that supplied. For physical-sequential access, no DSCB in the VTOC has a higher argument than that supplied. For a multiple DSCB request, the last DSCB in the VTOC was read and more DSCBs were requested.
8(08)	DEB invalid or omitted or not open to VTOC.	33(21)	EXTENTS keyword omitted, or supplied number of extents is zero.
9(09)	IOAREA=KEEP and user not authorized, or I/O area supplied and user not authorized	34(22)	CVAFDSM ACCESS=RLSEL format 0 DSCB already free (MAP=VTOC), or VIER already unallocated (MAP=INDEX) or volume space extent already unallocated (MAP=VOLUME).
10(0A)	Function not supported on OS VTOC.	42(2A)	VRF data supplied for write too long.
11(0B)	DSCB is not format-0 DSCB and VERIFY=YES.	43(2B)	Buffer list is for VIRs, but a DSCB buffer list is required.
12(0C)	MAPRCDS=YES and/or IXRCDS=KEEP but VTOC is nonindexed.	44(2C)	No buffer list entry found.
13(0D)	IXRCDS=KEEP not specified for CVAFDSM ACCESS=IXADD or IXDLT.		
14(0E)	CTAREA keyword omitted.		
15(0F)	UCB invalid, volume not mounted; VIO unit, not DASD.		
17(11)	DSCB length invalid for the function requested: 96 bytes for CVAFDIR ACCESS=WRITE,VERIFY=YES; 96 bytes for CVAFSEQ reading in data-set-name sequence; 140 bytes for CVAFSEQ reading in physical sequence.		
19(13)	UCB omitted and CVAF I/O area not supplied.		
22(16)	Data set name already supplied.		
23(17)	Invalid DSN supplied (44 X'FF' is a reserved data set name).		
24(18)	ARG keyword not supplied.		

Code	Meaning	Code	Meaning
45(2D)	Invalid DSCB buffer length (neither 96 nor 140) in buffer list entry, or VIR buffer length not equal to VIB VIR size.	132(84)	A level n vertical index entry pointer points to a VIER that is not at level n - 1.
46(2E)	Neither TTR nor CCHHR bits set in buffer list entry to be used in writing or reading a 140-byte DSCB.	133(85)	Level n horizontal index entry pointer points to VIER that is not at level n.
47(2F)	More than one of the TTR, CCHHR, and RBA bits set in the buffer list entry.	134(86)	Horizontal VIER/map pointer points to a VIR that is not a VIER/map (invalid ID in header).
48(30)	Both the DSCB and VIR bits set in the buffer list header.	135(87)	Horizontal map pointer points to VIR that is not one of the first n VTOC index records (n is recorded in VIXM field VIMRCDS), or the first record in the VTOC index is not a VIXM.
49(31)	RBA bit set in a buffer list entry for a DSCB buffer list.	136(88)	A level-1 index entry contains a CCHHR pointer that is outside the VTOC extent.
50(32)	TTR or CCHHR bit set in buffer list entry but buffer list header indicates buffer list is for a VIR.	137(89)	The first high-level VIER, as indicated in the VIB, does not have the flag bit set indicating it is the first high-level VIER. (This error is either recovered from by updating the VIB from the VIXM, or the error is changed to 129.)
52(34)	Combination of MAP and COUNT not supported.	138(8A)	The RBA of the VTOC index VIR does not match the RBA recorded in the header of the record.
53(35)	MAP omitted.	139(8B)	The first record of a map (VIXM, VPSM, or VMDS) is not one of the first n VTOC index records (n is recorded in the VIXM field, VIMRCDS).
54(36)	Buffer list for a VIR chained to or from a buffer list for a DSCB.	140(8C)	The data set name in a level n + 1 VIER entry is lower than the high key of the level n VIER that the level n + 1 VIER entry points to.
55(37)	Unauthorized caller and VIB not initialized.	141(8D)	First high-level VIER structure error bit is on.
56(38)	MAPRCDS=YES not specified but required.	142(8E)	I/O error indicating the VTOC index is not formatted correctly.
57(39)	Buffer list for a DSCB supplied, but buffer list for a VIR is required (in MAPRCDS or IXRCDS buffer list address in CVAF parameter list).	143(8F)	Either the index bit is zero, or the DOS bit is zero in the format-4 DSCB of a VTOC previously found to be an indexed VTOC.
58(3A)	Neither the VIR nor DSCB bit set in a buffer list header.	144(90)	No SYS1.VTOCIX.nnn data set name in a VTOC whose format-4 DSCB has the index bit on, indicating the VTOC has an index.
60(3C)	Invalid or conflicting setting of allocate option byte in space parameter		
127(7F)	I/O error occurred.		
128(80)	Reserved.		
129(81)	The first high-level VIER as indicated in the VIXM does not have the flag bit set indicating it is the first high-level VIER.		
130(82)	A horizontal or vertical VIER pointer is outside the RBA range of the VTOC index.		
131(83)	A vertical VIER pointer points to a VIR that is not a VIER (invalid ID in header).		

Code	Meaning	Code	Meaning
145(91)	The data set name in a level n + 1 VIER entry is higher than the high key of the level n VIER that the level n + 1 VIER entry points to	153(99)	The format-4 DSCB on an indexed VTOC is written with either the index- or DOS-bit zeroed on an indexed VTOC.
146(92)	Four or more high-level VIERS were encountered.	154(9A)	A space map extends over more than 10 VTOC index records.
147(93)	Too many levels in the VTOC index. The length of the search list was exceeded.	155(9B)	Data set name not found in section with key greater than or equal to the name being searched for. The VIER section containing the name is invalid.
148(94)	VIER invalid, because offset to last section is invalid.	156(9C)	Invalid VIER horizontal pointer. Horizontal pointer of VIER1 points to VIER2 whose high key is lower than or equal to the high key of VIER1.
149(95)	VIER invalid, because offset to last entry in a section is invalid.	157(9D)	Could not find entry in level-2 or higher VIER that matches the high key of the VIER.
150(96)	Media Manager initialization failed.	158(9E)	Invalid section length or invalid number of sections in a VIER header.
151(97)	Level-2 or higher VIER contains fewer than two entries.	159(9F)	The first high-level VIER pointed to by the VIB has an invalid ID in the header.
152(98)	RECOVER=YES specified, but the static text module (ICVIXST0) indicates recovery is not permitted.		

APPENDIX D. EXAMPLE OF AN OPEN INSTALLATION EXIT MODULE

PROCESSING IN IFGOEX0B

The following program listing is a sample of IFGOEX0B. The four subroutines (BUFNO, SCREEN, RLSE, and SQTY) show examples of the kind of processing that can be done in your installation's version of IFGOEX0B.

The BUFNO subroutine defaults the number of buffers for QSAM DCBs (DCBBUFNO) if the value is zero when the exit is given control. The block size in the DCB (DCBBLKSI) is used, together with a fixed amount of storage (64K bytes in the example) to determine a buffer number. A buffer number is limited to a fixed value (32 in the example). Storage quantity and maximum buffer number are contained in two tables, DAMAX and TPMAX, that are used for DASD devices and tape devices, respectively. Storage quantity is expressed in units of 1024 (1K) bytes. The values in the DAMAX and TPMAX tables can be altered by your installation.

The SCREEN subroutine determines those cases in which the succeeding subroutines, RLSE and SQTY, should be executed. DASD sequential and partitioned data sets being processed by BSAM or QSAM and opened for OUTPUT or OUTIN are selected. The VTOC data set and data sets starting with 'SYS1.' (system data sets) are excluded. An installation may want to make further selection tests.

REQUESTING PARTIAL RELEASE

The RLSE subroutine sets on the partial release indicators in the JFCB if the number of extents in the data set is less than a fixed value (8 in the example). It sets off the partial release indicators in the JFCB if the number of extents in the data set is equal or greater than a fixed value (8 in the example). Partitioned data sets are not processed, because they may be opened many times to write one new member for each OPEN/CLOSE.

UPDATING THE SECONDARY SPACE DATA

The SQTY subroutine provides a default secondary space quantity if none is specified. The default is one half of the primary space quantity if it is greater than one. If the primary quantity is zero, secondary is set to a fixed default number of tracks (5 in the example). If the primary quantity is one, secondary is set to the same fixed default (5); note that, in this case, the secondary quantity is in units of tracks, cylinders, or average blocks, depending on the unit of the primary quantity.

If the secondary space quantity is not zero, the SQTY subroutine tests the number of extents in the data set. If the number of extents is equal to or greater than a fixed value (10 in the example), then the secondary quantity is increased by 50% if it is greater than 1. It is set to a default quantity (5 in the example) if the secondary quantity is one; note that, in this case, the secondary quantity is in units of tracks, cylinders, or average blocks, depending on that of the primary quantity.

IFGOEX0B CSECT

XX

* * * * *

* FUNCTION = * * * * *

* FOUR SAMPLE ROUTINES ARE SUPPLIED. * * * * *

* * * * *

* BUFNO - DEFAULT DCBBUFNO * * * * *

* DCBBUFNO (NUMBER OF BUFFERS) IS DEFAULTED FOR * * * * *

* OPENS TO PHYSICAL SEQUENTIAL AND PARTITIONED DATA SETS * * * * *

* ON DASD AND TAPE USING QSAM, FOR WHICH DCBBUFNO IS ZERO. * * * * *

* DCBBUFNO FOR SYSIN, SYSOUT, TERMINAL, AND DUMMY DATA SETS * * * * *

* IS SET TO THE EQUATE, INOUTBNO, OR THE VALUE IN THE * * * * *

* FULLWORD, INOUTBN. * * * * *

* * * * *

* DCBBUFNO IS SET TO THE NUMBER OF DCBBLKSZ BUFFERS WHICH * * * * *

* FIT IN A GIVEN AMOUNT OF STORAGE. THE AMOUNT OF STORAGE IS * * * * *

* DEFINED BY THE EQUATES, DAMXK AND TPMXK (OR THE FULLWORDS * * * * *

* AT LABELS, DAMAXK AND TPMAXK), FOR DASD AND * * * * *

* TAPE, RESPECTIVELY. THE EQUATES DEFINE THE AMOUNT OF * * * * *

* STORAGE FOR BUFFERS IN UNITS OF 1024 (IF DAMXK IS 32, THEN * * * * *

* THE AMOUNT OF STORAGE IS 32K, OR 32768). * * * * *

* DAMXK OR TPMXK TIMES 1024 IS DIVIDED BY DCBBLKSI TO * * * * *

* DETERMINE THE NUMBER OF BUFFERS TO DEFAULT. * * * * *

* * * * *

* THE EQUATES, DAMXBNO AND TPMXBNO, OR THE FULLWORDS * * * * *

* AT LABELS, DAMAXBNO AND TPMAXBNO, * * * * *

* DEFINE THE MAXIMUM NUMBER OF BUFFERS TO BE * * * * *

* DEFAULTED FOR DASD AND TAPE IF THE CALCULATION, ABOVE, * * * * *

* RESULTS IN A LARGER NUMBER. * * * * *

* * * * *

* SCREEN - SCREEN OUT CASES FOR RLSE, SQTY * * * * *

* * * * *

* RLSE - SET OR ZERO PARTIAL RELEASE * * * * *

* THIS ROUTINE SETS PARTIAL RELEASE FOR DASD PS (NOT PO) DATA * * * * *

* SETS BEING OPENED FOR OUTPUT OR OUTIN. * * * * *

* * * * *

* PARTIAL RELEASE IS SET ON IF THE NUMBER OF EXTENTS IS LESS * * * * *

* THAN A QUANTITY DEFINED BY THE EQUATE, RLSE1, OR THE BYTE, * * * * *

* EXTRLSE1. * * * * *

* * * * *

* PARTIAL RELEASE IS SET OFF IF THE NUMBER OF EXTENTS IS NOT * * * * *

* LESS THAN A QUANTITY DEFINED BY THE EQUATE, RLSE0, OR THE * * * * *

* BYTE, EXTRLSE0. * * * * *

* * * * *

* SQTY - SET OR UPDATE SECONDARY SPACE QUANTITY * * * * *

* THIS ROUTINE UPDATES THE SECONDARY SPACE * * * * *

* QUANTITY FOR DASD PS OR PO DATA SETS BEING * * * * *

* OPENED FOR OUTPUT OR OUTIN. * * * * *

* * * * *

* IF THE SECONDARY QUANTITY IS NOT ZERO, * * * * *

* AND IF THE NUMBER OF EXTENTS IN THE DATA SET IS * * * * *

* AT LEAST EQUAL TO THE QUANTITY IN THE EQUATE, EXTSQT (OR * * * * *

* THE BYTE AT LABEL, EXTSQTY), THEN: * * * * *

* 1. IF THE SECONDARY QUANTITY IS GREATER THAN ONE, * * * * *

* SECONDARY QUANTITY IS INCREASED BY ONE HALF * * * * *

* (50%). * * * * *

* 2. IF THE SECONDARY QUANTITY IS ONE, * * * * *

* SECONDARY QUANTITY IS SET TO THE VALUE IN THE FULLWORD * * * * *

* AT LABEL, SQTYDFLT (EQUAL TO THE EQUATE, SQTYDFL). * * * * *

* * * * *

* IF THE SECONDARY QUANTITY IS NOT ZERO, * * * * *

* AND IF THE NUMBER OF EXTENTS IN THE DATA SET IS * * * * *

* LESS THAN THE QUANTITY IN THE EQUATE, EXTSQT (OR * * * * *

* THE BYTE AT LABEL, EXTSQTY), SECONDARY QUANTITY * * * * *

* IS LEFT UNCHANGED. * * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

* * * * *

```

*
* IF SECONDARY QUANTITY IS ZERO, IT IS SET TO ONE HALF
* OF PRIMARY QUANTITY IF PRIMARY IS NOT ZERO OR ONE.
* IF PRIMARY QUANTITY IS ZERO, THE SPACE TYPE IS SET TO TRACKS,
* AND SECONDARY QUANTITY IS SET TO THE VALUE IN THE FULLWORD
* AT LABEL SQTDFLT (EQUAL TO THE EQUATE, SQTDFL).
* IF PRIMARY QUANTITY IS ONE, SECONDARY QUANTITY IS SET TO
* VALUE IN THE FULLWORD AT LABEL SQTDFLT (EQUAL TO THE
* EQUATE, SQTDFL).
*
* NOTES = SEE BELOW
*
* DEPENDENCIES =
* CLASS ONE CHARACTER CODE. THE EBCDIC CHARACTER CODE
* WAS USED FOR ASSEMBLY. THE MODULE MUST BE REASSEMBLED
* IF A DIFFERENT CHARACTER SET IS USED FOR EXECUTION.
*
* RESTRICTIONS = NONE
*
* REGISTER CONVENTIONS =
* R1 OIEXL ADDRESS
* R2 DCB ADDRESS
* R3 UCB ADDRESS
* R4 DCB BLOCK SIZE
* R5 ADDRESS OF TPMAX OR DAMAX TABLES
* R6 EVEN REGISTER OF EVEN/ODD PAIR
* R7 ODD REGISTER OF EVEN/ODD PAIR
* R8 TIOT ENTRY ADDRESS
* R8 JFCB ADDRESS
* R10 FORMAT 1 DSCB ADDRESS
* R11 SAVE RETURN CODE
* R13 SAVE AREA ADDRESS
* R14 RETURN ADDRESS
* R15 BASE REGISTER
*
* PATCH LABEL = PATCH
*
* MODULE TYPE = CONTROL (OPEN, CLOSE, EOVS DATA MANAGEMENT)
*
* PROCESSOR = ASSEMBLER XF
*
* MODULE SIZE = SEE EXTERNAL SYMBOL DICTIONARY
*
* ATTRIBUTES = REENTRANT, REFRESHABLE, READ-ONLY, ENABLED,
* PRIVILEGED, SUPERVISOR STATE, KEY ZERO,
* LINK PACK AREA RESIDENT/PAGEABLE
*
* ENTRY POINT = IFGOEX0B
*
* PURPOSE = SEE FUNCTION
*
* LINKAGE =
* FROM IFG0196L:
* BALR 14,15
*
* INPUT = STANDARD LINKAGE CONVENTIONS
*
* OUTPUT = DCBBUFNO DEFAULTED
* PARTIAL RELEASE SET OR RESET
* CONTIGUOUS FLAG SET TO ZERO
* SECONDARY SPACE REQUEST MODIFIED
* RETURN CODE IN REGISTER 15
* 0 IF JFCB NOT MODIFIED
* 4 IF JFCB MODIFIED
*

```

```

*
* EXIT-NORMAL =
*       BR 14
*
* EXIT-ERROR =
*       NONE
*
* EXTERNAL REFERENCES = SEE BELOW
*
* ROUTINES = NONE
*
* DATA AREAS = NONE
*
* CONTROL BLOCK = NONE
*
* TABLES = NONE
*
* MACROS = MODESET, IECDIEXL, DCBD, IEFUCBOB, IEFTIOT1, IEFJFCBN,
*       IECSDSL1

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

REGISTER EQUATES

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
R1      EQU 1          OIEXL PARAMETER LIST ADDRESS
RDCB    EQU 2          DCB ADDRESS
RUCB    EQU 3          UCB ADDRESS
RBKSIZ  EQU 4          DCB BLOCK SIZE
RMAX    EQU 5          ADDRESS OF TPMAX OR DAMAX
REVEN   EQU 6          EVEN REGISTER OF EVEN/ODD PAIR
RODD    EQU 7          ODD REGISTER OF EVEN/ODD PAIR. HAS *
                        DCBBUFNO DEFAULT
RTIOT   EQU 8          TIOT ENTRY ADDRESS
RJFCB   EQU 9          JFCB ADDRESS
RDSCB   EQU 10         FORMAT 1 DSCB ADDRESS
RINCODE EQU 11         INTERNAL RETURN CODE
R12     EQU 12
RSAVE   EQU 13         SAVE AREA ADDRESS
RET     EQU 14         RETURN ADDRESS
RCODE   EQU 15         BASE REGISTER/RETURN CODE ON EXIT
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

RETURN CODE

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
MODJFCB EQU 4          RETURN CODE IF JFCB MODIFIED
        USING IFGOEX0B,RCODE
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

START OF SAMPLE PROGRAM

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
      B      AFTRID1
      DC     C'IFGOEX0B JDM1137 &SYSDATE'
      DC     C'IFGOEX0B JDM1137 05/01/81'
+ AFTRID1  SAVE (14,12)  SAVE REGISTERS
+AFTRID1  DS      0H
+          STM      14,12,12(13)  SAVE REGISTERS
          XR      RINCODE,RINCODE  ZERO RETURN CODE
          USING   OIEXL,R1  PARAMETER LIST
          BAL     RET,BUFNO  DEFAULT BUFNO
          BAL     RET,SCREEN  SCREEN OUT CASES WHERE RLSE, *
                                AND SQT SHOULD NOT BE CALLED
          BAL     RET,RLSE   SET PARTIAL RELEASE
          BAL     RET,SQTY   SET SECONDARY QUANTITY
EXIT     EQU     *          RETURN TO CALLER
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```

```

*****
* RETURN TO CALLER
*****
LR RCODE,RINCODE
RETURN (14,12),RC=(15) RESTORE REGISTER
+ L 14,12(13,0) RESTORE REGISTER 14
+ LM 0,12,20(13) RESTORE THE REGISTERS
+ BR 14 RETURN
BUFNO EQU * DEFAULT DCB BUFNO
*****
*
* DEFINE DEFAULT VALUES
* DAMXK = NUMBER OF K (1024) OF BUFFERS FOR DASD
* TPMXK = NUMBER OF K (1024) OF BUFFERS FOR TAPE
* DAMXBNO = MAXIMUM NUMBER OF BUFFERS FOR DASD
* TPMXBNO = MAXIMUM NUMBER OF BUFFERS FOR TAPE
* NOTE THAT DAMXBNO AND TPMXBNO MUST NOT BE GREATER THAN 255
*
*****
DAMXK EQU 64 64K BUFFERS FOR DASD
TPMXK EQU 64 64K BUFFERS FOR TAPE
DAMXBNO EQU 32 32 BUFFERS MAXIMUM FOR DASD
TPMXBNO EQU 32 32 BUFFERS MAXIMUM FOR TAPE
INOUTBNO EQU 1 DCBBUFNO DEFAULT FOR SYSIN, SYSOUT, *
AND DD DUMMY
ONEK EQU 10 SHIFT ARGUMENT TO MULTIPLY BY 1024
B AFTRID2
DC CL8'BUFNO' BUFNO ROUTINE ID
AFTRID2 BCR 0,RET NOP RETURN
L RDCB,OIEXPDCB PROTECTED COPY OF DCB
USING DCBD,RDCB
*****
* DO NOT PROCESS EXCP, BSAM, DSORG NOT PS OR PO,
* DCBBUFNO SPECIFIED
*****
TM DCBMACF1,DCBMRECP EXCP DCB?
BO RETBUFNO RETURN IF EXCP
TM DCBMACF1,DCBMRRD READ MACRO
BO RETBUFNO RETURN IF READ-NOT QSAM
TM DCBMACF2,DCBMRWRT WRITE MACRO
BO RETBUFNO RETURN IF WRITE-NOT QSAM
TM DCBDSRG1,DCBDSGPS+DCBDSGPO PS OR PO
BZ RETBUFNO EXIT IF NOT PS OR PO
CLI DCBBUFNO,0 IS DCBBUFNO SPECIFIED
BNE RETBUFNO RETURN IF DCBBUFNO SPECIFIED
*****
* DEFAULT DCBBUFNO TO 1 FOR SYSIN, SYSOUT, TERMINAL, DUMMY
*****
L RTIOT,OIEXTIOT TIOT ENTRY ADDRESS
USING TIOENTRY,RTIOT
L RODD,INOUTBN BUFNO DEFAULT FOR SYSIN/SYSOUT/ *
DD DUMMY
TM TIOELINK,TIOESSDS+TIOOTERM SYSIN/SYSOUT OR TERMINAL
BNZ STORE BRANCH IF SYSIN OR SYSOUT OR TERMINAL
L RJFCB,OIEXJFCB JFCB ADDRESS
USING INFMJFCB,RJFCB
CLC JFCBDSNM(L'NULLFILE),NULLFILE DUMMY DATA SET
BE STORE BRANCH IF DUMMY
*****
* EXIT IF NO UCB ADDRESS OR BLOCK SIZE NOT POSITIVE
*****
L RUCB,OIEXUCB UCB ADDRESS
LTR RUCB,RUCB ANY UCB?
BZ RETBUFNO EXIT IF NO UCB
LH RBKSIZ,DCBBLKSI DCB BLOCK SIZE
LTR RBKSIZ,RBKSIZ ANY BLOCK SIZE?
BNP RETBUFNO RETURN IF NO BLOCK SIZE

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*      GET TAPE OR DASD MAX TABLE
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
      USING UCBOB,RUCB
      TM      UCBTBYT3,UCB3DACC  DASD UCB?
      LA      RMAX,DAMAX        MAX TABLE FOR DASD
      BO      CALC              BRANCH IF DASD
      TM      UCBTBYT3,UCB3TAPE TAPE UCB?
      LA      RMAX,TPMAX        MAX TABLE FOR TAPE
      BZ      RETBUFNO         RETURN IF NOT DASD OR TAPE
CALC   EQU      *              DEFAULT DCBBUFNO
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*      CALCULATE DEFAULT BUFFER NUMBER
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
      USING MAX,RMAX
      XR      REVEN,REVEN      ZERO EVEN REG
      L       RODD,MAXBUF      MAXIMUM STORAGE FOR BUFFERS
      SLL     RODD,ONEK        SHIFT TO MULTIPLY BY 1024
      DR      REVEN,RBKSIZ     DIVIDE MAS BUFFER SPACE BY BKSI
      C       RODD,MAXBNO      ARE THERE TOO MANY BUFFERS?
      BNH     STORE           USE CALCULATION IF NOT TOO LARGE
      L       RODD,MAXBNO      USE MAXIMUM NUMBER OF BUFFERS
STORE  EQU      *              DEFAULT DCBBUFNO FOR USER/COPY DCB
      STC     RODD,DCBBUFNO    PUT IN PROTECTED COPY OF DCB
      L       RDCB,OIEXUDCB    USER DCB
      XR      REVEN,REVEN      MODESET USES REG 6 = REVEN
      MODESET KEYADDR=OIEXUKEY,WORKREG=6 GET IN USER KEY
+* /* MACDATE Y-3 77277                      @ZA26071*/
+* /*
+      IC      6,OIEXUKEY      GET KEY FROM SAVE LOCATION
+      SPKA    0(6)           SET PSW KEY
      STC     RODD,DCBBUFNO    PUT IN USER DCB
      MODESET EXTKEY=ZERO     BACK TO KEY ZERO
+* /* MACDATE Y-3 77277                      @ZA26071*/
+* /*
+      SPKA    0(0)           SET PSW KEY
RETBUFNO EQU *              RETURN FROM BUFNO
      BR      RET             RETURN
INOUTBN DC      A(INOUTBNO)   SYSIN/SYSOUT/DUMMY BUFNO DEFAULT
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*      MAX TABLE FOR TAPE
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
      DS      0F
      DC      CL8'TPMAX'      TPMAX ID
TPMAX   DS      0F
TPMAXK  DC      A(TPMXK)      MAXIMUM SIZE FOR BUFFERS IN UNITS *
                                OF 1024
TPMAXBNO DC A(TPMXBNO)      MAXIMUM NUMBER OF BUFFERS
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*      MAX TABLE FOR DASD
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
      DS      0F
      DC      CL8'DAMAX'      DAMAX ID
DAMAX   DS      0F
DAMAXK  DC      A(DAMXK)      MAXIMUM SIZE FOR BUFFERS IN UNITS *
                                OF 1024
DAMAXBNO DC A(DAMXBNO)      MAXIMUM NUMBER OF BUFFERS

```

```

SCREEN EQU * SCREEN OUT CASES WHERE RLSE, *
AND SQTY SHOULD NOT EXECUTE
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
* DO NOT PROCESS IF
* SYSIN/SYSOUT/TERMINAL
* DD DUMMY
* USER ASKS JFCB NOT BE RE-WRITTEN
* SYSTEM DATA SET ('SYS1.XXX')
* NON-DASD UCB
* NOT A FORMAT 1 DSCB
* EXCP DCB
* DSORG IN DCB IS NEITHER PS NOR PO
* DSORG IN DSCB IS NEITHER PS NOR PO
* NEITHER PUT NOR WRITE MACRO CODED IN DCB
* OPEN FOR OTHER THAN OUTPUT OR OUTIN
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
AFTRID3 B AFTRID3
DC CL8'SCREEN' SCREEN ROUTINE ID
L RTIOT, OIEXTIOT TIOT ENTRY ADDRESS
TM TIOELINK, TIOESSDS+TIOOTERM SYSIN/SYSOUT OR TERMINAL
BNZ EXIT EXIT IF SYSIN OR SYSOUT OR TERMINAL
L RJFCB, OIEXJFCB JFCB ADDRESS
CLC JFCBDSNM(L'NULLFILE), NULLFILE DUMMY DATA SET
BE EXIT EXIT IF DUMMY
CLC SYS1, JFCBDSNM SYS1.XXX DATA SET
BE EXIT EXIT IF SYSTEM DATA SET
TM JFCBTSDM, JFCNWRIT DON'T MODIFY JFCB
BO EXIT EXIT IF YES
L RUCB, OIEXUCB UCB ADDRESS
LTR RUCB, RUCB ANY UCB?
BZ EXIT EXIT IF NO UCB
TM UCBTBYT3, UCB3DACC DASD UCB?
BNO EXIT EXIT IF NOT DASD
L RDSCB, OIEXDSCB FORMAT 1 DSCB ADDRESS
USING DS1FMTID, RDSCB
CLI DS1FMTID, C'1' IS THIS A FORMAT 1 DSCB
BNE EXIT EXIT IF NOT
L RDCB, OIEXPDCB PROTECTED DCB ADDRESS
TM DCBMACF1, DCBMRPDCB EXCP DCB?
BO EXIT EXIT IF EXCP
TM DCBDSRG1, DCBDSGPS+DCBDSGPO PS OR PO DCB
BZ EXIT EXIT IF NOT PS OR PO
NC DS1DSORG, DS1DSORG IS DSORG SPECIFIED
BZ TSTMACRF TRUST DCB IF NOT SPECIFIED
TM DS1DSORG, DS1DSGPS+DS1DSGPO IS DATA SET PS OR PO
BZ EXIT EXIT IF NOT PS OR PO
TSTMACRF EQU * TEST MACRF IN DCB
TM DCBMACF2, DCBMRPUT PUT MACRO
BO TSTOOPT TEST OPEN OPTION
TM DCBMACF2, DCBMRWRT WRITE MACRO
BZ EXIT EXIT IF NOT WRITE
TSTOOPT EQU * TEST OPEN OPTION
TM OIEXOOPT, OIEXOOUT OPEN FOR OUTPUT
BO SCREENOK BRANCH IF YES
TM OIEXOOPT, OIEXOIN OPEN FOR OUTIN
BNO EXIT EXIT IF NO
SCREENOK EQU *
BR RET RETURN TO CALL RLSE, SQTY

```

RLSE EQU * SET PARTIAL RELEASE

```

*
* DEFINE DEFAULT VALUES
* RLSE0 = NUMBER OF EXTENTS. IF THE DATA SET HAS THIS
* NUMBER OF EXTENTS OR MORE, THEN PARTIAL RELEASE
* WILL NOT BE ALLOWED.
* RLSE1 = NUMBER OF EXTENTS. IF THE DATA SET HAS LESS THAN
* THIS NUMBER OF EXTENTS, PARTIAL RELEASE IS
* REQUIRED.
*
* NOTE THAT RLSE0 MUST NOT BE GREATER THAN RLSE1
*
* SETTING RLSE0 TO 17 OR GREATER WILL CAUSE THIS ROUTINE TO
* NEVER PREVENT A REQUEST FOR PARTIAL RELEASE
*
* SETTING RLSE1 TO 0 WILL CAUSE THIS ROUTINE TO
* NEVER FORCE A REQUEST FOR PARTIAL RELEASE
*
  
```

```

RLSE0 EQU 8 SET RELEASE BIT TO ZERO IF NUMBER OF *
EXTENTS EQUAL OR GREATER THAN THIS
RLSE1 EQU 8 SET RELEASE BIT TO ONE IF NUMBER OF *
EXTENTS LESS THAN THIS
B AFTRID4
DC CL8'RLSE' RLSE ROUTINE ID
AFTRID4 BCR 0,RET NOP RETURN
L RDSCB,OIEXDSCB FORMAT 1 DSCB ADDRESS
TM DS1DSORG,DS1DSGPO IS DATA SET PARTITIONED
BO TSTRLSE DO NOT SET RELEASE FOR PARTITIONED
CLC DSINOEPV,EXTRLSE1 FEW ENOUGH TO SET RELEASE
BNL TSTRLSE BRANCH IF NOT
L RJFCB,OIEXJFCB
OI JFCBIND1,JFCRLSE SET RELEASE
LA RINCODE,MODJFCB JFCB MODIFIED
B RETRLSE RETURN
TSTRLSE CLC DSINOEPV,EXTRLSE0 ENOUGH TO ZERO RELEASE
BL RETRLSE BRANCH IF NO
NI JFCBIND1,255-JFCRLSE ZERO RELEASE
LA RINCODE,MODJFCB JFCB MODIFIED
RETRLSE EQU * RETURN FROM RLSE
BR RET RETURN
DC CL8'RLSECONS' RLSE CONSTANTS ID
DS 0H
EXTRLSE1 DC AL1(RLSE1) IF FEWER THAN THIS NUMBER OF EXTENTS,*
PARTIAL RELEASE WILL BE SET
EXTRLSE0 DC AL1(RLSE0) IF THIS NUMBER OR MORE EXTENTS, *
PARTIAL RELEASE WILL BE ZEROED
  
```

```

SQTY    EQU    *           SET SECONDARY QUANTITY
*****
*
*      DEFINE DEFAULT VALUES
*      SQTYDFL = DEFAULT SECONDARY QUANTITY. THIS QUANTITY IS
*              SET IF THE SECONDARY QUANTITY IS ZERO AND THE
*              PRIMARY QUANTITY IS ZERO OR ONE. IT IS USED
*              IF SECONDARY QUANTITY IS ONE, AND THE NUMBER OF
*              EXTENTS IS EQUAL OR GREATER TO EXTSQT.
*      EXTSQT  = NUMBER OF EXTENTS. IF THE DATA SET HAS THIS MANY
*              EXTENTS OR MORE, THEN INCREASE SECONDARY QUANTITY.
*
*****
SQTYDFL EQU    5           DEFAULT SECONDARY QUANTITY
EXTSQT  EQU    10        IF DATA SET HAS THIS MANY EXTENTS, *
                          THEN INCREASE SECONDARY QUANTITY

AFTRID6 B      AFTRID6
DC      CL8'SQTY'       SQTY ROUTINE ID
BCR     0,RET           NOP RETURN
L       RJFCB,OIEXJFCB  JFCB ADDRESS
NC      JFCBSQTY,JFCBSQTY ANY SECONDARY QUANTITY
BZ      TSTPRIM        TEST PRIMARY IF NOT
L       RDSCB,OIEXDSCB  FORMAT 1 DSCB ADDRESS
CLC     DS1NOEPV,EXTSQT ENOUGH TO ADD TO SECONDARY QTY
BL      RETSQTY        BRANCH IF NOT
XR      RODD,RODD
ICM     RODD,7,JFCBSQTY GET SECONDARY QUANTITY
LR      REVEN,RODD     SAVE IN REVEN
SRL     REVEN,1        HALVE SECONDARY QUANTITY
LTR     REVEN,REVEN    IS SECONDARY ONE
BZ      SETDFLT        DEFAULT SECONDARY IF ONE
AR      RODD,REVEN    150% OF SECONDARY
B       STSQTY
TSTPRIM EQU    *       SECONDARY QUANTITY IS ZERO
NC      JFCBPQTY,JFCBPQTY IS PRIMARY QUANTITY ZERO
BZ      DFLTSQTY      DEFAULT SECONDARY
XR      RODD,RODD
ICM     RODD,7,JFCBPQTY
SRL     RODD,1        HALVE PRIMARY
LTR     RODD,RODD     IS PRIMARY ONE
BNZ     STSQTY        BRANCH IF NOT
SETDFLT EQU    *       USE QUANTITY IN SQTYDFLT
L       RODD,SQTYDFLT DEFAULT SECONDARY
B       STSQTY        STORE SECONDARY
DFLTSQTY EQU    *     PRIMARY AND SECONDARY ZERO
L       RODD,SQTYDFLT GET DEFAULT SECONDARY
TM      JFCBCTRI,JFCBSPAC
BNZ     STSQTY
CLI     DS1EXT1,X'01'  TRACK EXTENT
BE      DFLTTRK       YES -- SET TRACKS
CLI     DS1EXT1,X'81'  CYL EXTENT
BNE     RETSQTY       NO -- RETURN
OI      JFCBCTRI,JFCBCYL SET CYLINDER UNITS
B       STSQTY
DFLTTRK EQU    *     SET TRACK UNITS
OI      JFCBCTRI,JFCBTRK MAKE TRACK REQUEST
STSQTY EQU    *     STORE SECONDARY QTY
STCM    RODD,7,JFCBSQTY
LA      RINCODE,MODJFCB JFCB MODIFIED
RETSQTY EQU    *     RETURN FROM SQTY
BR      RET          RETURN
DS      OF
DC      CL8'SQTYCONS'  SQTY ROUTINE CONSTANTS ID
SQTYDFLT DC    A(SQTYDFL) DEFAULT SECONDARY QUANTITY
DC      AL1(0)        NOTE ONE BYTE OF ZERO BEFORE EXTSQTY
EXTSQT  DC    AL1(EXTSQT) IF DATA SET HAS THIS MANY EXTENTS, *
                          THEN ADD TO SECONDARY QUANTITY

```

```

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*          CONSTANTS / PATCH AREA
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
NULLFILE DC C'NULLFILE ' DD DUMMY DATA SET NAME
SYS1     DC C'SYS1.'      START OF SYSTEM DATA SET NAMES
        DS OF
PATCH   DC C'IFG0EX0B PATCH AREA '
        DC XL50'00'
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*          MAX TABLE MAPPING DSECT (MAPS TPMAX OR DAMAX)
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
MAX      DSECT
MAXBUF   DS A              MAXIMUM SIZE FOR BUFFERS
MAXBNO   DS A              MAXIMUM NUMBER OF BUFFERS
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*          DCB OPEN INSTALLATION EXIT PARAMETER LIST
*          - THE IECDIEXL MACRO IS IN SYS1.MACLIB
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
IECDIEXL
***** THE MACRO EXPANSION IS NOT SHOWN
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*          DCB - THE DCBD MACRO IS IN SYS1.MACLIB
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
DCBD DSORG=PS,DEV=DA
***** THE MACRO EXPANSION IS NOT SHOWN
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*          UCB - THE IEFUCBOB MACRO IS IN SYS1.AMODGEN
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
UCB     DSECT
        IEFUCBOB LIST=YES
***** THE MACRO EXPANSION IS NOT SHOWN
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*          TIOT - THE IEFTIOT1 MACRO IS IN SYS1.AMODGEN
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
TIOT    DSECT
        IEFTIOT1
***** THE MACRO EXPANSION IS NOT SHOWN
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*          JFCB - THE IEFJFCBN MACRO IS IN SYS1.AMODGEN
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
JFCB    DSECT
        IEFJFCBN LIST=YES
***** THE MACRO EXPANSION IS NOT SHOWN
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
*
*          FORMAT 1 DSCB - THE IECDLSL1 MACRO IS IN SYS1.AMODGEN
*
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
F1DSCB  DSECT
        IECDLSL1 (1)
***** THE MACRO EXPANSION IS NOT SHOWN
END

```

INDEX

A

ABE appendage 48-49
abnormal-end appendage
 See ABE appendage
access method routines, functions
 performed in I/O operations 36
accessing VTOCs and VTOC indexes 25-34
alias name
 of UCS images for JES2 174
allocate routine 4
altering DADSM processing 96-100
alternate track
 assigning 61, 62
AM operand
 in DEBCHK macro 126
APF (authorized program facility) 13
appendages
 abnormal-end (ABE) 48-49
 channel-end (CHE) 47
 end-of-extent (EOE) 46
 entry points 44
 listing in SYS1.PARMLIB 49
 naming convention 49
 page fix 72
 PCI 45
 programming restrictions 44
 register usage 44
 returns 44
 start-I/O (SIO) 45
assigning alternate track 61-62
associated programs
 DADSM 96
 IEHLIST 35
ATLAS macro
 coding example 62
 how to use 62
 operations performed 63
 return codes 63
 specification 61-62
 with track overflow option 61
authorized appendage list 49-50

B

BALANCE operand (TRKCALC macro) 135,
 137, 139
BFALN operand (DCB macro) 55
BFTEK operand (DCB macro) 55
bit maps
 of allocated cylinders and tracks 10
 of allocated DSCBs 11
 of allocated VIRS 11
block ID
 setting in JFCB for high-speed
 positioning 123
block multiplexor programming notes 51
BSAM (basic sequential access method)
 defaulting buffer number 103
BUFCB operand (DCB macro) 55
buffer
 releasing 32
buffer lists

format
 of entries 28-30
 of header 28
function 28
how created 28
releasing 32
BUFL operand (DCB macro) 55
BUFNO
 defaulting in OPEN installation exit
 example (QSAM) 225
BUFNO operand (DCB macro) 55

C

catalogs
 dummy module 176
CCW (channel command word) 41
 See also channel program
CENDA operand (DCB macro) 53
channel program
 appendages used with 43
 execution 41-42
 initiation 41-42
 related 43
 restrictions or modification 42
 translation, virtual addresses to
 real addresses 73-74
channel-end appendage
 See CHE appendage
CHE appendage 47
checking the DEB 124-128
checkpoint data set
 processed with EXCP macro 56
CLOSE macro
 used with EXCP macro 66
 used with XDAP macro 79
codes
 returned with error message 222-224
 routing and descriptor 221
command retry 50
communication vector table (CVT) mapping
 macro 111
completion codes 70, 80
 See also return codes
 following use of EXCP macro 70
 following use of XDAP macro 80
control blocks
 DCB 40, 51
 DEB 40
 ECB 40, 69
 FCB 156
 general description 40
 IOB 40, 67-69
 PIRL 43, 131
control password 89
conversion
 actual track address to relative
 track address 82
 of sector value for RPS devices 82
 relative track address to actual
 track address 80
copy operation
 requirements 14
creating protected data sets 86
CVAF

- serialization 26
- CVAF processing, of GTF trace 35
- CVAF, volume identification to 26
- CVAFDIR macro
 - examples 196-205
 - how to use 30
 - return codes 183
 - syntax 178-183
 - uses 178
- CVAFDSM macro
 - example 217-220
 - how to use 34
 - return codes 189
 - syntax 185-189
 - uses 185
- CVAFSEQ macro
 - examples 206-217
 - how to use 32
 - return codes 194
 - syntax 190
 - uses 190
- CVAFTST macro
 - example 217-220
 - return codes 195
 - syntax 195
 - uses 195
- CVFCTN field of CVPL
 - contents 28
 - definitions 28
- CVPL (CVAF parameter list)
 - format 27
 - function 27
 - when created 27
- CVSTAT codes 222
- CVT (communication vector table) mapping
 - macro 111

D

- DADSM
 - routines 1
- DADSM pre/postprocessing exits 96
 - format-1 DSCB passed by IGGPRE00 99
 - parameter list (IEPL) 98
 - register contents 100
 - return codes from IGGPRE00 100
 - system control block addresses 99
- DASD (direct access storage devices)
 - reading and writing to
 - XDAP macro 75-78
- data extent block (DEB)
 - See DEB
- data management ABEND installation
 - exit 107-109
 - parameter list (OAIXL) 108
 - register contents 109
 - return codes 109
- data set control block (DSCB)
 - See DSCB
- data set security
 - See password protection
- DCB fields used with EXCP macro 51-58
- DCB OPEN exit
 - return codes 105
- DCB OPEN installation exit 100-105
 - example
 - defaulting buffer number (QSAM) 225

- requesting partial release 225
- updating secondary space data 225
- parameter list (OIEXL) 102
- register contents 104
 - when executed 101-102
- DCBDIRECT field of DCB 54
- DCBFDAD field, maintaining 54
- DCBIFLGS field of DCB, permanent I/O
 - error indicators 42
- DCBOFLGS field of DCB, meanings of bit
 - settings 65-66
- DCBTRBAL field, maintaining 56
- DD operand (TRKCALC macro) 136, 137, 139
- DDNAME operand (DCB macro) 52
- DDR (dynamic device reconfiguration),
 - repositioning tape data sets 52
- DEB (data extent block)
 - fields 70
 - obtaining 26
 - use with EXCP macro 40
 - validating 124-128
- DEBCHK
 - register contents 127
- DEBCHK macro
 - functions of 124-128
 - return codes 128
 - specification 125-128
- defaulting buffer number, for BSAM 103
- defaulting buffer number, for QSAM
 - in OPEN installation exit
 - example 225
- defaulting the DCB buffer number 103
- defective track
 - See assigning alternate track
- define extent command 41
- deleting a data set
 - coding example 20
 - macro instructions for 18
 - when volume not mounted 19
- DEMOUNT parameter (MSGDISP macro) 148
- DEN operand (DCB macro) 57
- DEQ at demount facility 121
- DEV operand (DCB macro) 55-57
- device characteristics 112-117
- device-dependent parameters in
 - DCB 55-58
- DEVTAB operand (TRKCALC macro) 135, 137
- DEVTYPE macro
 - for RPS devices 113
 - output from 112-117
 - return codes 115
 - specification 112-113
- DSCB (data set control block)
 - format-1 not found 105
 - general information 1
 - reading from VTOC by actual device
 - address
 - coding example 18
 - macro specifications 17
 - return codes 18
 - reading from VTOC by data set name
 - coding example 16
 - macro specifications 15
 - return codes 17
- DSECT expansions
 - See CVT, IEFJFCBN, IEFUCBOB, TRKCALC
- DSN order, initiating 32
- DSORG operand (DCB macro) 54, 56-57

E

ECB fields
 with EXCP macro 69
 with XDAP macro 79
 end-of-extent appendage
 See EOE (end-of-extent) appendage
 end-of-volume
 See EOVS macro
 EODAD operand (DCB macro) 54
 EOE appendage 46
 EOE operand (DCB macro) 52
 EOVS macro
 and Format-1 DSCB not found 105
 with EXCP macro 65-66
 with XDAP macro 79
 error handling 34
 error messages
 See messages, CVAFF VTOC error
 error recovery
 from system or user errors 34
 procedures 42
 event control block (ECB) fields
 with EXCP macro 69
 with XDAP macro 79
 examples
 of CVAFFDIR macro 196-205
 of CVAFFSEQ macro 206-217
 of CVAFFST and CVAFFDSM
 macros 217-220
 of OPEN installation exit module 225
 EXCP macro
 control blocks used with
 DCB 51-58
 DEB 70
 ECB 69
 IOB 67-69
 in problem programs 38
 in real storage 70
 in system control programs 37
 in V=R address space 39
 macro specification 60
 macros used with
 ATLAS 61
 CLOSE 66
 EOVS 65-66
 OPEN 58-60
 multivolume data set requirement 60
 EXCPVRS macro 70-71
 executing channel program
 in problem programs 38
 in real storage 70
 in system control programs 37
 exit routine
 See DADSM pre/postprocessing exits
 See data management ABEND
 installation exit
 See DCB OPEN installation exit
 See format-1 DSCB
 EXLST operand (DCB macro) 54
 expiration date
 overriding 20
 EXTEND operand (OPEN macro) 58, 121
 EXTEND operand (OPEN macro) The DADSM
 extend routine obtains 4

F

FCB (forms control buffer) image
 adding image to SYS1.IMAGELIB 156
 adding to SYS1.IMAGELIB 169
 JES2 Support 174
 retrieving from SYS1.IMAGELIB 173
 rules 156
 fixing data areas with EXCPVRS 71
 format
 of buffer list entry 29-30
 of buffer list header 28
 of CVPL (CVAFF parameter list) 27
 of DADSM pre/postprocessing exit
 parameter list 98
 of OIEXL (OPEN installation exit
 parameter list) 102
 of VIER 8
 of VIER index entries 8
 of VTOC maps 12
 format 0-6 DSCB 1
 format-1 DSCB
 not found installation exit
 (IFGOEXOA) 105-107
 register contents 106
 return codes 107
 reading from VTOC 15
 format-1 DSCB not found exit
 register contents 106
 return codes 107
 format-1 DSCB not found installation
 exit
 parameter list 106
 forms control buffer (FCB)
 See FCB image
 foundation block of DCB 52
 FUNCTN operand (TRKCALC macro) 134-139

G

GEN parameter (MSGDISP macro) 152
 GTF trace of CVAFF processing
 See CVPL (CVAFF parameter list)

H

high-speed 3480 positioning 123

I

I/O appendages
 See appendages
 I/O devices
 characteristics 112
 IDAL (indirect address list) 74
 IDAL (indirect data address list) 73
 IDAW programming notes
 IEAAPPOO, authorized appendage list 49
 IEBUPDTE program
 SYS1.PARMLIB 49-50
 use in listing appendages in 49

IEPCNVTV (relative track address to actual track address conversion routine) 80
 IECPLTV (actual track address to relative track address conversion routine) 82
 IECOSCR1 (sector conversion routine) 83
 IEFJFCBN macro 111
 IEFUCBOB macro 110
 IEHATLAS program 62
 IEPL (DADSM pre/postprocessing exit parameter list) 98
 IFGOEX0A program
 See format-1 DSCB
 IFGOEX0B program
 See DCB OPEN installation exit
 IFG0199I program
 See data management ABEND installation exit
 IGGPOSTO (DADSM postprocessing exit)
 See DADSM pre/postprocessing exits
 IGGPRE00 (DADSM preprocessing exit)
 See DADSM pre/postprocessing exits
 IGGPRE00 program
 IGGUCSIT macro 166
 IMSK operand (DCB macro) 54
 indexed access, initiating 32
 indexed VTOCs, compared to nonindexed VTOCs 12
 indexing feature for 3211 174
 indirect address list (IDAL) 74
 indirect data address list (IDAL) 73
 initializing DASD volumes 1
 interruption handling procedures 42
 IOB chain modification 132
 IOB fields
 with EXCP macro 67-69
 with XDAP macro 79
 IOBAD operand (DCB macro) 54
 IOBSENS fields with ATLAS macro 62

J

JES2
 Printer support 174, 175
 JFCB (job file control block) 111, 117, 118, 120, 122, 124
 See also RDJFCB macro
 macros used with
 OPEN 122
 RDJFCB 118-120
 mapping macro 111
 modifying 103, 118-124
 processing 117-124
 job file control block (JFCB)
 See JFCB

K

K operand (TRKCALC macro) 136, 137, 139
 KEYLEN operand (DCB macro) 57

L

LABEL operand (DD statement)
 password protected data set 85, 86
 library character set modules 156
 LIST operand
 CVT macro 111
 IEFJFCBN macro 111
 IEFUCBOB macro 111
 locate record command 41

M

MACRFE=(E) operand (DCB macro) 52
 macros, data management
 See also CVAFDIR macro
 See also CVAFDSM macro
 See also CVAFSEQ macro
 See also CVAFTST macro
 ATLAS 61
 CLOSE
 used with EXCP macro 66
 used with XDAP macro 79
 CVAF (VTOC access)
 syntax 195
 uses 25
 CVT 111
 DCB 51, 58
 DEBCHK 125-128
 DEVTYPE 112-117
 EOVB
 and Format-1 DSCB not found 105
 used with EXCP macro 65-66
 used with XDAP macro 79
 EXCP 60
 EXCPVR 70-71
 IEFJFCBN 111
 IEFUCBOB 110
 MSGDISP 143-155
 OBTAIN 15-18
 OPEN
 and Format-1 DSCB not found 105
 for JFCB 120, 122
 used with EXCP macro 58-60
 PROTECT 88-95
 PURGE 128-132
 RDJFCB 117-120
 REALLOC 140-142, 143
 RENAME 22-25
 RESTORE 128-129, 133
 SCRATCH 18
 TRKCALC 133-139
 used with XDAP macro 75
 XDAP 75-78
 maintaining 95
 See also PROTECT macro
 PASSWORD data set 88-95
 VTOC (volume table of contents) 15-25
 mapping macros
 CVT 111
 IEFJFCBN 111
 IEFUCBOB 110
 TRKCALC 138
 maps of allocated space
 for cylinders and tracks 10
 for DSCBs 11
 for VIRs 11

MAXSIZE operand (TRKCALC macro) 136, 137
 messages, CVAF VTOC error associated codes
 descriptor codes 221
 return codes 222-224
 routing codes 221
 text and explanation 221
 MF operand
 DEBCHK macro 127
 TRKCALC macro 136-137
 MODE operand (DCB macro) 57
 modifying
 channel program during execution 42
 IOB chain 132
 JFCB
 in OPEN installation exit 103, 118-124
 MOUNT parameter (MSGDISP macro) 144
 MSGDISP macro 143
 displaying a Demount message 148
 displaying a Mount message 144
 displaying a Ready message 147
 displaying a Verify message 146
 Providing the Full Range of Display Options 152
 Resetting the Message Display 151
 return codes 154
 multivolume data set
 processing with EXCP macro 59

N

name
 VTOC index 6
 nonindexed VTOCs, compared to indexed 12
 nonpageable address space, V=V 41
 NOPWREAD protection-mode indicator 89
 NOWRITE protection-mode indicator 89

O

OAIXL (data management ABEND installation exit parameter list) 108
 OBTAIN macro 15-17
 obtaining a sector number (RPS devices) 83
 OIEXL (OPEN installation exit parameter list) 102
 OPEN installation exit
 See DCB OPEN installation exit
 OPEN macro
 and DEQ at demount facility 121
 and Format-1 DSCB not found 105
 getting control from 102
 TYPE=J
 example 59
 invoking 121
 specification 123
 used with EXCP macro
 dummy data set restriction 58
 label processing 58
 procedures performed 58
 volume disposition 58
 used with XDAP macro 76
 open processing 100
 after IFGOEXOB gets control 101

before IFGOEXOB gets control 101
 opening a VTOC, restriction on changing contents 119
 OPENJ (OPEN, TYPE=J) 122
 operational requirements 14
 OPTCD=Z operand (DCB macro) 53
 OUTINX operand (OPEN macro) 54, 121
 output data set
 maintaining DCBBLKCT field 52

P

page boundary 73
 page fix
 appendage 72
 list 72
 pageable address space, V=R 39
 parameter list
 See CVPL (CVAF parameter list)
 See DADSM pre/postprocessing exits
 See data management ABEND installation exit
 See DCB OPEN installation exit
 See format-1 DSCB not found installation exit
 partial release via JFCB
 modification 104
 in OPEN installation exit example 225
 password
 See also PROTECT macro
 control 89
 counter maintenance 88
 data set concatenation 88
 parameter list 91
 ADD a record 91
 DELETE a record 93
 LIST a record 94
 REPLACE a record 92
 protection mode indicator 89
 record 86
 protection mode indicator 86
 secondary 89
 standard label restriction 84
 terminating protection 87
 volume switching 87
 PASSWORD data set
 characteristics 86
 creating 86
 password protecting data sets 84-95
 password protection
 for VTOC indexes 14
 PCI (program controlled interruption) appendage 45
 PCI operand (DCB macro) 52
 PCIA operand (DCB macro) 52
 pg=end.DADSM pre/postprocessing exits 100
 PGFX appendage 72
 physical sequential access, initiating 33
 PIRL (purged I/O restore list)
 use in restoring I/O requests 43, 131, 133
 posting completion code in ECB
 following use of EXCP macro 69
 following use of XDAP macro 79
 postprocessing
 See DADSM
 PREFIX operand (IEFUCBOB macro) 111
 preprocessing

See DADSM
 printer image
 universal character set (UCS) 156
 program controlled interruption (PCI)
 appendage 45
 programming notes 50
 PROTECT macro
 See also password
 parameter list 93
 ADD function 91
 DELETE function 93
 LIST function 94
 REPLACE function 92
 return codes 95
 syntax 90
 use with PASSWORD dataset 84-94
 protecting a VTOC index
 with passwords 13
 with RACF 13
 protection mode indicator 89
 PRTSP operand (DCB macro) 58
 PURGE macro
 adding to macro library 128
 definition 128
 parameter list 130, 132
 return codes 132
 specification 130
 purged I/O restore list 43, 130, 133
 PWRITE protection-mode indicator 89
 PWRITE protection-mode indicator 89

Q

QSAM (queued sequential access method)
 defaulting buffer number 103

R

R operand (TRKCALC macro) 136, 137, 139
 RACF (Resource Access Control Facility)
 renaming a data set 22
 scratching a data set 19
 use with VTOCs and VTOC indexes 13
 RDJFCB macro
 coding example 118
 exit list entry for 119
 invoking DEQ at demount 121
 return codes 120
 specification 118
 RDY parameter (MSGDISP macro) 147
 reading
 data from index maps 31
 DSCBs 30, 31, 32
 VIR 30-31
 reading and modifying a JFCB 117-120
 READPSWD module 85
 REALLOC macro 140-143
 return codes 142
 RECFM operand (DCB macro) 54
 recovering from errors 34
 recovering from permanent DASD I/O error
 See ATLAS macro
 register
 conventions for appendages 43
 usage by conversion routines 81, 82
 usage by I/O supervisor 43
 REGSAVE operand (TRKCALC macro) 136,
 137

related channel program 43
 related requests 43
 REMOVE operand (TRKCALC macro) 135, 137
 RENAME macro
 dummy module 176
 specification 22
 renaming a data set
 coding example 24
 macro specification 22
 with password protection 25
 REPOS operand (DCB macro) 52
 requesting partial release via JFCB
 modification 104
 requirements
 for APF 13
 for copy, restore, operations 14
 RESET parameter (MSGDISP macro) 151
 Resource Access Control Facility
 See RACF (Resource Access Control
 Facility)
 restore chain modification 133
 RESTORE macro
 adding to macro library 128
 definition 128
 specification 133
 restore operations, requirements 14
 restoring IOBs 133
 restrictions
 when scratching, renaming,
 allocating 12
 return codes
 ATLAS macro 63
 CVAFDIR macro 183
 CVAFDSM macro 189
 CVAFSEQ macro 194
 CVAFTST macro 195
 DEBCHK macro 128
 DEVTYPE macro 115
 from DADSM exits 100
 from DADSM pre/postprocessing
 exits 100
 from data management ABEND exit 109
 from DCB OPEN exit 105
 from format-1 DSCB not found
 exit 107
 IECPCNVT 82
 MSGDISP macro 154
 OBTAIN macro 17
 RDJFCB macro 120
 REALLOC macro 142
 TRKCALC macro 139
 with error message 222-224
 RKDD operand (TRKCALC macro) 136, 137
 routine
 See exit routine
 RPS (rotational position sensing)
 devices
 used with XDAP macro 82

S

SCRATCH macro
 coding example 20
 general description 18
 scratching a data set
 when volume not mounted 19
 secondary passwords 89
 secondary space data, updating in OPEN
 installation exit 104
 example 225
 sector, address in XDAP macro 78, 82

- securing a VTOC
 - with passwords 13
- securing a VTOC index
 - See also protecting a VTOC index with RACF 13
- seek 41
- serialization, CVAF 26
- SIO appendage
 - description 45
 - for EXCPVR 72
- SIOA operand (DCB macro) 52
- space map
 - of allocated cylinders and tracks 10
 - of allocated DSCBs 10
 - of allocated VIRs 10
- SSCH (start subchannel) instruction
 - description 45
- STACK operand (DCB macro) 57
- stand-alone seek 41
- standard label restriction, password data sets 84
- start subchannel (SSCH) instruction
 - See SSCH
- start-I/O appendage
 - description 45
 - for EXCPVR 72
- system control blocks 99
 - mapping macros for
 - CVT 111
 - JFCB (IEFJFCBN) 111
 - UCB (IEFUCBOB) 111
- system macro instructions 110
- SYS1.IMAGELIB data set
 - adding a UCS image to 157
 - maintaining 156
 - UCS images in 157

T

- tape volumes
 - DEQ at demount facility 121
- testing for a VTOC index 195
- trace of CVAF processing 35
- track
 - assigning alternate 61
 - calculating capacity 133-139
- translation of channel program
 - by I/O supervisor
 - in V=R address space 72-73
 - in V=V address space 41
 - in your own program 72-73
- TRKBAL operand (TRKCALC macro) 134, 137, 138
- TRKCALC macro 133-139
 - return codes 139
- TRKCAP operand (TRKCALC macro) 134, 139
- TRTCH operand (DCB macro) 57
- TYPE operand
 - DEBCHK macro 125
 - TRKCALC macro 135, 137, 139

U

- UCB (unit control block)
 - mapping macro 113
- UCB operand
 - TRKCALC macro 135, 137
- UCS (universal character set) image
 - adding to SYS1.IMAGELIB 157
 - adding to the UCS image table 162
 - for JES2 174
 - verifying 162
- UCS image table
 - adding aliases 162
 - adding image names 162
 - contents 163
 - entry format 162
 - modifying entries 166
 - structure 162
- unit check with ATLAS macro 62
- unit control block (UCB)
 - See UCB (unit control block)
 - getting information from
 - See DEVTYPE macro
- universal character set (UCS)
 - See UCS (universal character set) image
- updating secondary space data in OPEN installation exit 104
- user exit routine
 - See exit routine

V

- V=R address space, EXCP operations
 - in 39
- V=V address space 41
- validating the DEB 124-128
- VERIFY parameter (MSGDISP macro) 146
- VIER (VTOC index entry record)
 - characteristics 7
 - contents 7
 - format 8
 - function 6
 - how chained together 10
 - when created 9
- VIR (VTOC index record)
 - kinds 6
 - length 6
- virtual IDAWs
- VIXM (VTOC index map)
 - format 12
 - function 11
- VMDS (VTOC map of DSCBs)
 - format 12
 - function 11
- volume
 - identification to CVAF 26
 - volume label 2
 - volume switching 59
 - volume table of contents
 - See VTOC (volume table of contents)
- VPSM (VTOC pack space map)
 - format 12
 - function 10
- VTOC (volume table of contents)
 - maintaining
 - description 1
 - index 15, 18
 - using OBTAIN macro 15-18

- using RENAME macro 22-25
- using SCRATCH macro 18-22
- VTOC access macros
 - See CVAFDIR macro
 - See CVAFDSM macro
 - See CVAFSEQ macro
 - See CVAFTST Macro
- VTOC index
 - contents 6
 - how to list 35
 - how to protect 14
 - name 6
 - relationship to VTOC 6
 - structure 6
 - testing for 195
- VTOC index entry record (VIER)
 - See VIER (VTOC index entry record)
- VTOC index map (VIXM)
 - See VIXM (VTOC index map)
- VTOC index record (VIR)
 - See VIR (VTOC index record)
- VTOC map of DSCBs (VMDS)
 - See VMDS (VTOC map of DSCBs)
- VTOC pack space map (VPSM)
 - See VPSM VTOC pack space map)

W

- WAIT macro
 - used with EXCP macro 38
- WRITE protection mode indicator 25
- writing DSCBs 31
- writing VIRs 31

X

- XDAP channel program 80
- XDAP macro
 - control blocks used with 75, 79
 - macros required with
 - CLOSE 79
 - EOV 79
 - OPEN 75-76
 - specification 77-78
- XENDA operand (DCB macro) 53

Numerics

- 1403 Printer
 - JES2 Support 174
- 31-bit IDAW
- 3203 printer
 - JES2 175
 - output from DEVTYPE 115
- 3211 Printer
 - indexing feature 174
 - JES2 Support 174
- 3262 Model 5 Printer
 - default FCB image 157
 - UCS image table 162
 - UCS image table for 165
- 3480 tape drive
 - output from DEVTYPE macro 116
- 3800 printer
 - output from DEVTYPE macro 115
- 4245 printer
 - default FCB image 157
 - output from DEVTYPE macro 115
 - UCS image table 162
- 4248 Printer
 - default FCB image 157
 - UCS image table 162

**Reader's
Comment
Form**

MVS/XA System-Data Administration
GC26-4010-2

This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate.

Note: Do not use this form to request IBM publications. If you do, your order will be delayed because publications are not stocked at the address printed on the reverse side. Instead, you should direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

If you have applied any technical newsletters (TNLs) to this book, please list them here: _____

Chapter/Section _____

Page No. _____

Comments:

Note: Staples can cause problems with automated mail sorting equipment. Please use pressure sensitive or other gummed tape to seal this form.

If you want a reply, please complete the following information.

Name _____ Phone No. (____) _____

Company _____

Address _____

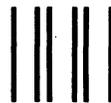
Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Fold and tape

Please do not staple

Fold and tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

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

POSTAGE WILL BE PAID BY ADDRESSEE

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



Fold and tape

Please do not staple

Fold and tape



**Reader's
Comment
Form**

MVS/XA System-Data Administration
GC26-4010-2

This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate.

Note: Do not use this form to request IBM publications. If you do, your order will be delayed because publications are not stocked at the address printed on the reverse side. Instead, you should direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

If you have applied any technical newsletters (TNLs) to this book, please list them here: _____

Chapter/Section _____

_____ Page No. _____

Comments:

Note: Staples can cause problems with automated mail sorting equipment. Please use pressure sensitive or other gummed tape to seal this form.

If you want a reply, please complete the following information.

Name _____ Phone No. (____) _____

Company _____

Address _____

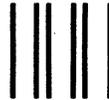
Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Fold and tape

Please do not staple

Fold and tape

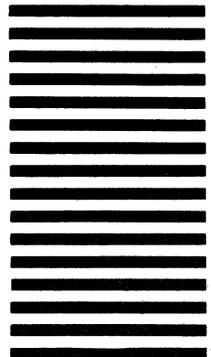


**NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES**

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

POSTAGE WILL BE PAID BY ADDRESSEE

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



Fold and tape

Please do not staple

Fold and tape





MVS/Extended Architecture
System-Data Administration

File Number S370-34

GC26-4010-02



Printed in U.S.A.