

SC34-0404-0

LICENSED
PROGRAM

File No. S1-32

IBM Series/1

**Event Driven Executive
Indexed Access Method
Version 2 Guide**

Program Numbers: 5719-AM4
5719-UT5

SC34-0404-0

LICENSED
PROGRAM

File No. S1-32

IBM Series/1

**Event Driven Executive
Indexed Access Method
Version 2 Guide**

Program Numbers: 5719-AM4
5719-UT5

First Edition (July 1981)

This edition applies to the Event Driven Executive Indexed Access Method Version 2 Modification Level 0: Program Number 5719-AM4, until otherwise indicated by new editions or technical newsletters.

Use this publication only for the purpose stated in the Preface.

Changes are periodically made to the information herein; any such changes will be reported in subsequent revisions or Technical Newsletters.

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

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

This publication could contain technical inaccuracies or typographical errors. A form for readers' comments is provided at the back of this publication. If the form has been removed, address your comments to IBM Corporation, Information Development, Department 27T, P. O. Box 1328, Boca Raton, Florida 33432. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

PREFACE

HOW TO USE THIS BOOK

The material in this section is a guide to using this book. It defines the purpose, audience, and content of the book as well as listing aids for using the book and background materials.

PURPOSE

The purpose of this publication is to describe how to use the Indexed Access Method Version 2 to set up indexed files and to develop application programs using indexed files.

AUDIENCE

This manual is intended for use by:

- Application designers to design applications which use the Indexed Access Method Version 2.
- Application programmers to develop applications which use the Indexed Access Method Version 2.

Applications for the Series/1 can be developed in several languages. Unless otherwise noted in this section, material in this book is intended for use in the development of applications in any of the following languages:

- COBOL
- EDL
- PL/I

HOW THIS BOOK IS ORGANIZED

This book describes the Indexed Access Method in the following order:

- Chapter 1, "Introduction" provides an overview of the Indexed Access Method.
- Chapter 2, "Using the Indexed Access Method," provides a brief description of what indexed files are, how to set up an indexed file, and application program request statements.

- Chapter 3, "Defining Primary Index Files," describes the format of the primary index file and how to use the \$IAMUT1 utility to set up your indexed files.
- Chapter 4, "Loading the Primary Index File," describes loading data records into a primary index file using an application program.
- Chapter 5, "Building a Secondary Index," provides information on using secondary keys, what a secondary index is and does, and how to set up and load a secondary index.
- Chapter 6, "Processing the Indexed File," describes how to process the indexed file with an application program.
- Chapter 7, "Coding The Indexed Access Method Requests," provides information needed to code EDL applications which use the Indexed Access Method. This chapter is intended only for EDL application developers.
- Chapter 8, "The \$IAMUT1 Utility," provides information needed to use \$IAMUT1, including the completion codes it generates.
- Chapter 9, "The \$VERIFY Utility," provides information needed to use \$VERIFY.
- Chapter 10, "Storage and Performance Considerations" describes the storage and performance characteristics of the Indexed Access Method and how to tailor the Indexed Access Method to the processing requirements of your installation.
- Chapter 11, "Error Recovery" describe some of the error recovery procedures available for use with Indexed Access Method applications.
- Chapter 12, "Installing the Indexed Access Method," provides an overview of the installation process.
- Appendix A, "Summary of Calculations," provides a summary of calculations for calculating the various blocks which make up indexed files.
- Appendix B, "Preparing Indexed Access Method Programs," provides an overview of preparing an Indexed Access Method application and a sample \$JOBUTIL procedure for an EDL application.
- Appendix C, "Coding Examples," provides comprehensive examples of Indexed Access Method programs. This appendix is for application developers using EDL, COBOL, or PL/I as their application programming language.

AIDS IN USING THIS PUBLICATION

Illustrations in this book are enclosed in boxes. Many illustrations display screens generated while using the Event Driven Executive system. In

those cases where the actual data exceeds the size of the box, the information may be illustrated in a modified format.

In display screens appearing in this manual, operator input is shown in bold type. This highlighting is for illustrative purposes only, to distinguish data entered by the operator from that generated by the system.

Several other aids are provided to assist you in using this book:

- A bibliography
 - Lists the books in the Event Driven Executive library along with a brief description of each book
 - Lists related publications and materials
- A glossary defines terms and abbreviations
- A common index includes entries from each book in the Event Driven Executive library

RELATED PUBLICATIONS

Related publications are listed in the Bibliography.

SUBMITTING AN APAR

If you have a problem with the Series/1 Event Driven Executive services, you are encouraged to fill out an authorized program analysis report (APAR) form as described in the IBM Series/1 Authorized Program Analysis Report (APAR) User's Guide, GC34-0099.

CONTENTS

Chapter 1. Introduction	IG-1
What The Indexed Access Method Does	IG-1
Indexed Access Method Features	IG-1
Devices Supported	IG-3
Languages Compatible With Indexed Access Method	IG-4
Components of Indexed Access Method	IG-4
Chapter 2. Using the Indexed Access Method	IG-5
Your Data Record	IG-5
Setting Up An Indexed File Using \$IAMUT1	IG-6
Processing The Indexed File	IG-13
Summary	IG-14
Chapter 3. Defining Primary Index Files	IG-17
Primary Index Files	IG-17
Data Record Primary Key	IG-17
Random and Clustered Record Inserts	IG-18
Defining The File Structure With \$IAMUT1	IG-19
Designing Indexed Files Using \$IAMUT1 - Option 1	IG-21
Option 1	IG-21
Designing Indexed Files Using \$IAMUT1 - Option 2	IG-24
Option 2	IG-24
Indexed Access Method Blocks	IG-24
Data Blocks	IG-26
Free Space	IG-27
Index Blocks	IG-30
Primary Index Blocks (PIXB)	IG-30
Second-level Index Blocks (SIXB)	IG-35
Higher-level Index Block (HIXB)	IG-37
Free Pool	IG-39
File Control Block	IG-40
File Structure Types	IG-40
Option 2 Examples	IG-42
Example 1: Allocating Free Records	IG-43
Example 2: Allocating Free Records and Free Blocks	IG-45
Example 3: Allocating Reserved Data Blocks	IG-48
Example 4: Allocating Reserved Index Entries	IG-51
Example 5 - Defining a Totally Dynamic File	IG-54
Designing Indexed Files Using \$IAMUT1 - Option 3	IG-57
\$IAMUT1 - Option 3	IG-58
Defining, Creating, and Loading a File - Summary	IG-59
Chapter 4. Loading The Primary Index File	IG-61
Loading the Primary Index File	IG-61
Loading Base Records using \$IAMUT1	IG-62
Loading Base Records From An Application Program	IG-65
Loading Base Records From a Sequential File in Random Order	IG-65

Chapter 5. Building a Secondary Index	IG-67
Secondary Keys	IG-67
The Directory	IG-67
Allocating and Inserting Entries in a Directory	IG-69
Secondary Index	IG-74
Defining and Loading A Secondary Index	IG-75
Example 1: Defining A Secondary Index Using \$IAMUT1	IG-77
Option 1	IG-77
Option 2	IG-79
Option 3	IG-82
Loading a Secondary File With an Application Program	IG-84
Chapter 6. Processing The Indexed File	IG-87
Connecting and Disconnecting the Indexed File	IG-87
Connecting	IG-87
Disconnecting	IG-88
Accessing the Indexed File	IG-89
Direct Reading	IG-89
Direct Updating	IG-90
Sequential Reading	IG-90
Sequential Updating	IG-91
Inserting Records	IG-93
Deleting Records	IG-93
Extracting Indexed File Information	IG-93
Maintaining the Indexed File	IG-94
File Backup and Recovery	IG-94
Recovery Without Backup	IG-95
Reorganizing an Indexed File	IG-95
Dumping an Indexed File	IG-96
Deleting an Indexed File	IG-96
Verifying an Indexed File	IG-96
Chapter 7. Coding the Indexed Access Method Requests	IG-97
Request Functions Overview	IG-97
Coding Indexed Access Method Requests	IG-99
CALL Function Descriptions	IG-101
DELETE - Delete Record	IG-101
DISCONN - Close File	IG-104
ENDSEQ - End Sequential Processing	IG-106
EXTRACT - Get File Information	IG-108
GET - Get Record	IG-111
GETSEQ - Get Record (Sequential Mode)	IG-114
LOAD - Open File for Record Loading	IG-118
PROCESS - Open File	IG-122
PUT - Put Record into File	IG-127
PUTDE - Delete Previously Read Record	IG-130
PUTUP - Update Record	IG-133
RELEASE - Release Record	IG-136
EDL CALL Functions Syntax Summary	IG-138
Indexed Access Method Return Codes Summary	IG-139
Chapter 8. The \$IAMUT1 Utility	IG-141
\$IAMUT1	IG-142
\$IAMUT1 Commands	IG-142

BF - Tailor the Indexed Access Method Buffers	IG-144
DF - Define Indexed File	IG-146
DI - Display Parameter Values	IG-149
DR - Invoke Secondary Index Directory Functions	IG-150
AL - Allocate Directory	IG-151
DE - Delete Directory Entry	IG-153
EN - End Directory Function	IG-154
IE - Insert Entry	IG-155
LE - List Entries	IG-156
UE - Update Directory Entry	IG-158
EC - Control Echo Mode	IG-160
EF - Display Existing Indexed File Characteristics	IG-161
LO - Load Indexed File	IG-162
NP - Deactivate Paging	IG-166
PG - Select Paging	IG-167
PP - Define Paging Partitions	IG-168
PS - Get Paging Statistics	IG-169
RE - Reset Parameters	IG-170
RO - Reorganize Indexed File	IG-171
SE - Set Parameters	IG-173
UN - Unload Indexed File	IG-184
\$IAMUT1 Completion Codes	IG-186
Chapter 9. The \$VERIFY Utility	IG-189
\$VERIFY Functions	IG-189
Invoking \$VERIFY	IG-190
\$VERIFY Input	IG-190
Invoking \$VERIFY From a Terminal	IG-191
Invoking \$VERIFY From a Program	IG-191
\$VERIFY Example	IG-192
FCB Report	IG-193
FCB Extension Report	IG-196
Free Space Report	IG-197
\$VERIFY Messages	IG-200
File Error Messages	IG-200
Error recovery procedure	IG-200
Other Messages	IG-201
\$VERIFY Storage Requirements	IG-201
Using Default Working Storage Requirements	IG-201
Modifying Working Storage Requirements	IG-201
Summary	IG-202
Chapter 10. Storage and Performance Considerations	IG-203
Determining Storage Requirements	IG-203
The Indexed Access Method Packages	IG-203
Indexed Access Method Storage Environment	IG-204
Performance	IG-205
Data Paging	IG-206
Other Performance Considerations	IG-209
Chapter 11. Error Recovery	IG-213
Handling Errors	IG-213
Error Exit Facilities	IG-213
System Function Return Codes	IG-215

The Data-Set-Shut-Down Condition	IG-215
\$ILOG - Error Logging Facility	IG-216
Deadlocks and the Long-Lock-Time Condition	IG-218
Verifying Requests and Files	IG-219
Chapter 12. Installing the Indexed Access Method	IG-221
Installation Procedures	IG-221
Installing The Indexed Access Method	IG-221
Assembling And Executing The Installation Verification Program	IG-222
Appendix A. Summary of Calculations	IG-223
Appendix B. Preparing Indexed Access Method Programs	IG-229
A Sample \$JOBUTIL Procedure and Link-Edit Control Data Set	IG-230
Appendix C. Coding Examples	IG-233
EDL Indexed Access Method Coding Example	IG-233
EDL Indexed Access Method Coding Example	IG-234
COBOL Indexed Access Method Coding Example	IG-238
PL/I Indexed Access Method Coding Example	IG-245
Bibliography	IG-255
Event Driven Executive Library	IG-255
Summary of Library	IG-257
Other Event Driven Executive Programming Publications	IG-260
Other Programming Publications	IG-261
Series/1 System Library Publications	IG-261
Glossary of Terms and Abbreviations	IG-263
Common Index	IG-277

FIGURES

1. Indexed File Logical Structure	IG-25
2. Data Block Format Example	IG-27
3. Cluster Example	IG-31
4. High-level Index Structure	IG-38
5. Indexed File with Free Records	IG-44
6. Indexed File with Free Records/Blocks	IG-47
7. Indexed File with Reserved Data Blocks	IG-50
8. Indexed File with Reserved Index Entries	IG-53
9. Totally Dynamic Indexed File	IG-56
10. Loading and Inserting Records	IG-62
11. Protocol for Sequential Updating	IG-92
12. Indexed Access Method Storage Environment	IG-205
13. Plot of Data Paging Area Sizes	IG-208
14. Volume Space Requirements	IG-221

CHAPTER 1. INTRODUCTION

The Indexed Access Method licensed program is a data management facility that executes on an IBM Series/1 processor under the Event Driven Executive Supervisor and Emulator, Version 3.1 or later. The Indexed Access Method provides keyed access to each of your individual data records.

WHAT THE INDEXED ACCESS METHOD DOES

This licensed program builds, maintains, and accesses a data structure called an indexed file

Your data records can be loaded by the Indexed Access Method utility, \$IAMUT1, or they can be loaded using an application program. Data records can then be added, deleted, modified, or accessed quickly and efficiently for processing by your application program. When reorganization of an indexed file is required the utility can be used to unload and reorganize the file.

When this licensed program is used, each of your records is identified by the contents of a predefined field called a key. The Indexed Access Method builds and maintains an index for those keys and through this index fast access to each record is provided. Your data records can be accessed either by key, or sequentially in ascending key sequence, using Indexed Access Method requests.

INDEXED ACCESS METHOD FEATURES

The Indexed Access Method offers the following features:

- Record access by a primary key or secondary keys - You can access records in an indexed file by one or more keys. Secondary keys use a separate index and Indexed Access Method provides the connection between the primary index files and secondary indexes. Duplication of secondary key fields is permitted.
- Support for high insert and delete activity - Free space can be distributed throughout the file and in a free-pool at the end of the file so that new records can be inserted. The space occupied by a deleted record is immediately available for inserting a new record.
- Direct and sequential access - You can access records either randomly by key, or sequentially in ascending key sequence.
- Data paging - You can improve Indexed Access Method performance by using data paging. With this feature active, the Indexed Access Method retains recently-used blocks of data records resident in main storage.

- **Dynamic file structure** - A dynamic file structure adjusts itself as needed to handle record additions and deletions. This provides a quick and easy method of designing an indexed file.
- **Concurrent access to a single file by several requests** - These requests can be from one or more programs. Data integrity is maintained by a file-, block-, and record-level locking system that prevents other programs from accessing the portion of the file being modified.
- **Implementation as a separate task** - A single copy of the Indexed Access Method executes and coordinates all requests. A buffer pool supports all requests and optimizes the space required for physical I/O; the only buffer required in an application program is the one for the record being processed.
- **Input records** - Either blocked or unblocked input records are accepted.
- **\$IAMUT1** - A utility program that allows you to maintain a secondary index directory, create, format, load, unload, and reorganize an indexed file. The load and unload functions accept either blocked or unblocked records.
- **\$VERIFY** - A utility program that allows you to check the integrity of the index structure, print control blocks, and print a free space report for an indexed file.
- **Error logging** - If multiple error return codes occur, errors are logged in the system error log.
- **\$ILOG** - The error log entries can be printed by using the \$ILOG utility.
- **File compatibility** - Files created by the Event Driven Executive Indexed Access Method are compatible with those created by the IBM Series/1 Realtime Programming System Indexed Access Method licensed program, 5719-AM1 and 5719-AM2 provided that the block size is a multiple of 256.
- **Data protection** - All input/output operations are performed by system functions. Therefore, all data protection facilities offered by the system also apply to indexed files. The following additional data protection is provided:
 - The exclusive option specifies that the file is for the exclusive use of a requester.
 - File-level, block-level, and record-level locking automatically prevents two requests from accessing the same file, the same block, or the same data record simultaneously.

- The immediate write back option causes all file modifications (delete, insert, update) to be written back to the file immediately.
- Accidental key modification for primary keys is prevented to help ensure that your index matches the corresponding data.
- Distribution packaging - The Indexed Access Method is distributed with the following variations available:
 - A full function package that is intended to be totally resident.
 - A full function package which uses an overlay structure.
 - A totally resident package without data paging.
 - A package without data paging which uses an overlay structure.

DEVICES SUPPORTED

The Indexed Access Method supports indexed files on the following direct access devices:

- 4962 Disk Storage Unit
- 4963 Disk Subsystem
- 4964 Diskette Unit
- 4965 Diskette Unit
- 4966 Diskette Magazine Unit

In addition, the Indexed Access Method supports IBM 4969 Magnetic Tape Unit for loading and unloading indexed files with the \$IAMUT1 utility.

Input control and output display devices supported for the \$IAMUT1, \$ILOG, and \$VERIFY utilities are:

- IBM 3101 Display Station
- IBM 4978 Display Station
- IBM 4979 Display Station
- IBM 4973, 4974, and 4975 Printers
- Teletypewriter Model 33/35 or ASCII equivalent

LANGUAGES COMPATIBLE WITH INDEXED ACCESS METHOD

The following programming languages can be used to code Indexed Access Method programs.

- COBOL
- EDL
- PL/I

COMPONENTS OF INDEXED ACCESS METHOD

The Indexed Access Method consists of the following components:

- Four load modules from which you can select to support your application program Indexed Access Method requests. These load modules are named:
 - \$IAM (full function with overlay)
 - \$IAMRS (full function resident)
 - \$IAMNP (overlay without data paging)
 - \$IAMRSNP (resident without paging)

The module you select will be named \$IAM after installation.

- A load module, \$IAMSTGM, which is used to obtain the data paging area, if the data paging feature is requested.
- A set of object modules that you may use to generate a customized load module. If you use one of the four supplied load modules, you do not need the object modules.
- The object module, IAM, is called a link module. You include IAM with your application program using the linkage editor to provide the interface to the Indexed Access Method.
- Two copy code modules for inclusion in EDL programs, IAMEQU and FCBEQU. IAMEQU provides symbolic parameter values for constructing CALL parameter lists. FCBEQU provides a map of the file control blocks.
- Load modules for each of the Indexed Access Method utilities \$IAMUT1, \$VERIFY, and \$ILOG.

CHAPTER 2. USING THE INDEXED ACCESS METHOD

The purpose of this chapter is to familiarize you with some fundamentals of the Indexed Access Method. Some of the features mentioned in the previous chapter will be described only in part here so that a basic example can be constructed. The purpose of this example is to demonstrate the ease with which you can establish an indexed file and to help you select which parts of the book apply directly to your application requirements.

YOUR DATA RECORD

The data records you wish to process with the Indexed Access Method have the following specific requirements:

- the records must contain a common field that can be used as a primary key
- each record must have a unique primary key
- the initial records to be loaded must be in ascending order by the primary key
- all records that make up an indexed file must be of the same length

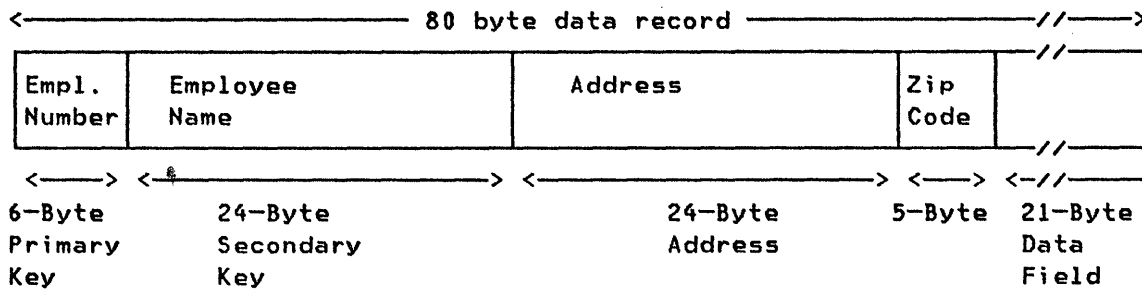
The primary key is any field you designate within your data records. The key field must begin at the same location in each record. Each key field must have the same length. The key in each record must be unique within the file (data set).

The data records that you will initially load must be in ascending order, based upon the field you use as the key. If your data records are not ready to be loaded when you define your primary indexed file, the records can be loaded later by an application program or with the LO (load) command of \$IAMUT1.

Your application might use an employee number as the primary key in an indexed file for some applications. You might want to define secondary keys, such as employee name, for the same file for other applications. Using secondary keys requires a secondary index to be defined. Defining a secondary index and using secondary keys is described in "Chapter 5. Building a Secondary Index."

Whether you use the \$IAMUT1 utility to load your data records into an indexed file from a sequential file, or load them with an application program, you must know the format of your input data record.

Following is a sample record layout. Although the primary key is shown starting in position 1, it could have been anywhere in the record.



The records used for our example have the following attributes:

- block size 256 bytes
- record size 80 bytes
- primary key length 6 bytes
- key position 1

SETTING UP AN INDEXED FILE USING \$IAMUT1

Use the Indexed Access Method utility program, \$IAMUT1, to set up an indexed file. After this utility is loaded into the system for execution, the utility displays a sequence of prompts. The prompts are questions displayed on a terminal one at a time to which you can reply using the terminal keyboard. Responding to the questions causes the utility to perform the required steps to:

1. set up the structure of the file (space for records to be loaded, free space for inserts, and an index)
2. allocate a data set (the utility prompts you for a data set and volume name and calls \$DISKUT3 to allocate space for the indexed file)
3. define and format the indexed file
4. load the data records into the indexed file

Loading and using the SE (set parameter) command of the \$IAMUT1 utility is described here for the purpose of our example, however, for a complete description of \$IAMUT1 see "Chapter 8. The \$IAMUT1 Utility."

The responses for our example are shown in **bold face type** inside the box. The bold bracketed numbers at the left, outside the box, identify explanatory remarks that we have written below the box using the same bracketed numbers. Of course these brackets and explanations do not appear on the screen when \$IAMUT1 is being used.

The \$IAMUT1 Indexed Access Method utility can be loaded with the Event Driven Executive operator command \$L \$IAMUT1.

When \$IAMUT1 is loaded the first prompt is displayed as follows:

```
[1] ENTER COMMAND (?): SE
```

[1] Entering the letters SE (set parameters), followed by pressing the ENTER key, causes four options to be displayed:

```
SET FILE DEFINITION PARAMETERS
0 = EXIT
[2] 1 = SIGNIFICANT PARAMETERS
    2 = ALL PARAMETERS
    3 = PARAMETERS FROM EXISTING INDEXED DATASET
ENTER OPTION: 1
```

[2] The response digit '1', causes prompts to follow which allows you to define an indexed file with a minimum of information. This response cause a one line prompt to be displayed.

Note: Although the following prompts are displayed one line at a time when using the utility, the prompts and responses are listed here in logical groups for simplicity in describing them.

[3]	SECONDARY INDEX (Y/N):? N		
		DEFAULT	NEW VALUE
[4]	RECORD SIZE		0:80
[5]	KEY SIZE		0:6
[6]	KEY POSITION		1:1
[7]	BLOCKING FACTOR (RECORDS PER BLOCK)		1:3
[8]	NUMBER OF BASE RECORDS		0:5
[9]	ESTIMATED TOTAL RECORDS		6:20
[10]	TYPE OF INSERT ACTIVITY(C=CLUSTERED,R=RANDOM)		C:R
[11]	DATA SET SIZE IN EDX RECORDS:	15	
[12]	INDEXED ACCESS METHOD RETURN CODE:	-1	
[13]	SYSTEM RETURN CODE:	-1	
[14]	CREATE/DEFINE FILE (Y/N):? Y		

[3] The first prompt, "SECONDARY INDEX (Y/N):?" asks if you are specifying a secondary index. The response was N for no, because we are defining the parameters for a primary indexed file.

[4] The second prompt, "RECORD SIZE" requests the length that the records are to be in the indexed file which you are defining.

Note that there are two columns near the right-hand edge of the display. The column on the left is headed by the word "DEFAULT". In the default column the values are listed that will be used in setting up the file if no value is supplied in the response (only the ENTER key is pressed). The column on the right, headed "NEW VALUE" is where the decimal value is placed from your keyboard response, followed by pressing the ENTER key.

In this example we are using a record length of 80.

[5] The "KEY SIZE" prompt is for the length of the primary key in the data record. In this example we are using a key which is 6 bytes long.

[6] Our key field begins in position 1 of the data record.

[7] We are requesting that our indexed file be blocked with 3 records in each 256-byte block.

[8] The number of base record slots to be defined is 5. This number is based on the number of data records we plan to load. You cannot load more records than this value, however, it does not restrict you from inserting new data records in the free (empty) slots later.

[9] The total number of records that we anticipate that this resultant indexed file will ever contain is 20.

[10] The type of record insert activity is to be R (random). The records added to this file will be inserted by an application program when those records are available.

The choice of random or clustered is based on the type of record additions that are anticipated. Random is chosen when the records to be added are expected to be evenly distributed throughout the file.

Clustered is chosen when the records to be added are expected to be in groups, relative to their range in key value.

[11] Following the previous response the system will display the number of records required to contain an indexed file using the parameters you have supplied.

[12] The Indexed Access Method return code (-1) indicates that the parameters you supplied are acceptable; no Indexed Access Method rules have been violated.

[13] The system return code (-1) should always be -1 if the Indexed Access Method return code is -1. If any errors are encountered, the return code may provide additional information.

[14] If you have verified that the parameters you entered are correct, the data set (file) size in EDX records is acceptable, and the return codes are both -1, you can reply Y and the file will be defined and created.

If you wish to change any of the parameter values that you previously supplied, respond N to this prompt and you will be prompted for the next command. To re-enter your responses, reply SE and the prompt sequence will be repeated.

A Y in response to this prompt causes the next prompt sequence to begin.

```

[15] ENTER DATA SET (NAME,VOLUME):  IAMFILE,EDX003
      NEW DATA SET IS ALLOCATED
[16] DO YOU WANT IMMEDIATE WRITE-BACK?  Y .
[17] INVOKE LOAD/REORGANIZE AFTER CURRENT FUNCTION (Y/N)?  Y
      DEFINE IN PROGRESS
      DATA SET SIZE IN EDX RECORDS:           15
      INDEXED ACCESS METHOD RETURN CODE:       -1
      SYSTEM RETURN CODE:                     -1
[18] INVOKE LOAD (L), REORGANIZE (R) OR END (E) ?L
      LOAD ACTIVE

```

[15] The data set and volume name you reply to this prompt is what \$DISKUT3 uses to allocate a data set for your file. A successful allocation results in the information message "NEW DATA SET IS ALLOCATED".

[16] The immediate write back option is recommended for most applications. It means that we want any record in the indexed file that we process with our application program to be written back to the indexed file immediately. Otherwise, the record will be held in a buffer until that buffer is needed by the Indexed Access Method.

[17] If you have data records in a sequential data set which you want \$IAMUT1 to load for you at this time, reply Y to this prompt.

If you are going to load the data records with your application program, reply N.

For our example we responded Y which causes the following information messages to be displayed: "DEFINE IN PROGRESS". The information messages then inform you with the data set (file) size in EDX records. The system return code and Indexed Access Method return code values at the completion of the file allocation and formatting are displayed for your information.

[18] Because of the Y response to the previous prompt, this prompt is generated to allow you to:

- load base records (as shown in this example)
- reorganize an existing indexed file for loading into the file being defined
- end the current SE command session

We are going to load records during this session so the response is Y. Following the 'LOAD ACTIVE' information message, the prompts continue.

```

[19] $FSEDIT FILE RECSIZE = 128
[20] INPUT RECORD ASSUMED TO BE      80 BYTES. OK?: N
[21] ENTER RECORD SIZE: 128
[22] ENTER INPUT  BLOCKSIZE (NULL = UNBLOCKED): 256
[23] ENTER INPUT  DATA SET (NAME,VOLUME): SEQ01,EDX003
      INPUT REC GT OUTPUT REC. TRUNCATION WILL OCCUR
[24] OK TO PROCEED:? Y
      LOAD IN PROCESS
      END OF INPUT DATA SET
[25] ANY MORE DATA TO BE LOADED?: N
      5 RECORDS LOADED
      LOAD SUCCESSFUL

```

The next sequence of prompts refers to the input data set containing the data records that are going to be loaded into the indexed file.

[19] The utility accepts input records which have been prepared by the Event Driven Executive utility \$FSEDIT. The \$FSEDIT record size is specified as 128.

[20] Because the output data set (indexed file) records are 80 bytes, this prompt determines whether the input sequential data set is also an 80 byte record data set.

If you use the Event Driven Executive edit utilities to prepare your data records for input to the Indexed Access Method, remember that these utilities place one 80-byte line from \$FSEDIT in a 128-byte record. The first record begins at location 1, and the second record begins at location 129. Two of these 128-byte records make one 256-byte EDX record.

Because we used \$FSEDIT, we responded N.

[21] This prompt requests the input data record attributes.

Because our input data records were created by \$FSEDIT, our 80-byte records were converted to 128-byte records. Therefore, our response is 128.

[22] The Indexed Access Method utility, \$IAMUT1 accepts your records as either unblocked (one record per block) or blocked (more than one record per block) input. The utility prompts you for the block size of the input data set being loaded. If the input data set is unblocked, reply to the block size prompt by pressing the Enter Key. See "Blocked and Unblocked Sequential Data Sets" on page IG-164 for a description of blocked and unblocked sequential data sets.

If your input data records are unblocked sequential, reply by pressing the Enter Key. If your input is blocked sequential, reply with the actual blocksize that was used to prepare your input data records.

Our example uses blocked sequential records, created on every line by \$FSEDIT, with a blocksize of 256.

[23] Reply to this prompt with your input data record data set and volume name. Our response was SEQ01,EDX003.

[24] This prompt verifies whether truncation of the input records is acceptable. Because our record size specified is actually 80 bytes long, but we responded 128 because \$FSEDIT converts the records to 128 bytes, the following warning message is displayed. "INPUT REC GT OUTPUT REC. TRUNCATION WILL OCCUR" This means that the extra bytes attached by \$FSEDIT to our 80-byte data records will now be removed. The response is Y.

The information message "LOAD IN PROCESS" tells us that \$IAMUT1 is reading the input data set and loading the input data records into the base record slots. The information message "END OF INPUT DATA SET" indicates that the end-of-file condition, on the input data set, has been encountered.

[25] This prompt allows you to specify another input data set, if more data records are to be loaded from another data set. In this example, only 1 data set is being used and the response of N caused the records loaded statistics to be displayed, followed by the "LOAD SUCCESSFUL" message.

The design of an indexed file varies according to your application. A comprehensive approach to designing your indexed files begins with "Defining The File Structure With \$IAMUT1" on page IG-19.

PROCESSING THE INDEXED FILE

Now that the indexed file has been defined, formatted, and loaded with data records, the file is ready for an application program to access any of the records in the indexed file for processing. An application program might use the following EDL coded requests to open the indexed file and retrieve a record.

```

*
* OPEN THE INDEXED FILE FOR PROCESSING
*
      .
      .
[1]      CALL  IAM,(PROCESS),IACB,(DS1),(OPENTAB),(SHARE)
*
* PERFORM A DIRECT RETRIEVAL OF THE RECORD WHOSE KEY IS JONES PW
*
[2]      CALL  IAM,(GET),IACB,(BUFF),(KEY1)
      .
      .
KEY1     TEXT  'JONES PW'
OPENTAB  DATA F'0'
          DATA A(IAMERR)
          DATA F'0'
IACB     DATA F'0'
```

[1] This Indexed Access Method request opens the primary index file in process mode so that other requests can be issued for processing records in the indexed file.

[2] This Indexed Access Method request retrieves a record from the indexed file. The primary key of this record contains the name 'JONES PW'.

Functions of the Requests

Following is a list of functions that you can perform using the Indexed Access Method requests in your application program:

Initiate general purpose access to an indexed file with a PROCESS request. After the PROCESS request has been issued, any of the following functions can be requested:

- Direct reading - Retrieving a single record independently of any previous request.
- Direct updating - Retrieving a single record for update; complete the update by either replacing or deleting the record.

- Sequential reading - Retrieving the next logical record relative to the previous sequential request.

The first sequential request can access the first record in the file or any other record in the file.

- Sequential updating - Retrieving the next logical record for update; complete the update by either replacing or deleting the record.
- Inserting - Placing a single record, in its logical key sequence, into the indexed file.
- Deleting - Removing a single record from the indexed file.
- Extracting - Extracting data that describes the file.

Note that the update functions require more than one request.

When a function is complete, another function may be requested, except that a sequential processing function can be followed only by another sequential function. You can terminate sequential processing at any time by issuing a DISCONN or ENDSEQ request. An end-of-data condition also terminates sequential processing.

A complete list of the Indexed Access Method requests, the operand descriptions, and correct syntax is described in "Chapter 7. Coding the Indexed Access Method Requests" on page IG-97. There are also coding examples using the Indexed Access Method requests in three programming languages in "Appendix C. Coding Examples." The languages used in the examples are Event Driven Language, COBOL, and PL/I. The purpose of these examples is not to show any particular application, but to help you when planning and writing your application program.

SUMMARY

This chapter has introduced some fundamentals of using the Indexed Access Method. The references in this chapter to other chapters in this manual were placed there to help you select the specific information you need for your application. A list of those references is repeated here to assist you in locating the detailed information on the listed subjects.

- For a complete description of \$IAMUT1 see "Chapter 8. The \$IAMUT1 Utility"
- A comprehensive approach to designing your indexed files is described in "Chapter 3. Defining Primary Index Files"
- Defining a secondary index for using secondary keys is described in "Chapter 5. Building a Secondary Index"
- Description of blocked and unblocked sequential data sets is described in "Blocked and Unblocked Sequential Data Sets" on page IG-164

- The complete list of Indexed Access Method requests, the operand descriptions, and correct syntax is described in "Chapter 7. Coding the Indexed Access Method Requests"
- Guide line information on processing the indexed file is located in "Chapter 7. Coding the Indexed Access Method Requests." This guideline information should be read prior to planning and coding your application program.

CHAPTER 3. DEFINING PRIMARY INDEX FILES

This chapter presents the following major topics:

- Primary Indexed Files
- Designing Indexed Files Using \$IAMUT1 - option 1
- Designing Indexed Files Using \$IAMUT1 - option 2
 - Indexed Access Method Blocks
 - Index Blocks
 - File Control Block
 - File Structure Types
 - Option 2 Examples
- Designing Indexed Files Using \$IAMUT1 - option 3
- Defining/Creating, and Loading A File - Summary

This chapter provides information for defining indexed files and is arranged according to your option selection when using \$IAMUT1. The beginning of the chapter has information which applies to any type of primary index file design. That general information section is followed immediately with an example using \$IAMUT1, option 1. The option 2 section is next and contains information that you will need to know prior to designing an index file with \$IAMUT1, option 2. The fourth section applies to using \$IAMUT1, option 3.

PRIMARY INDEX FILES

A primary index file contains data records, a multilevel index, control information, and it can optionally contain free space.

Free space can be distributed throughout the file and at the end of the file. Free space provides areas for inserting new records and is described later.

In an indexed file, the records are arranged in ascending order by key.

DATA RECORD PRIMARY KEY

The primary key can be any field within your data record that you select, however, it must meet the following requirements:

- the selected field must start at the same location in each record
- all portions of the key field must be contiguous
- the primary key length cannot exceed 254 bytes
- the field must contain data that is unique within the data set

Defining the Key

Define a single key field by specifying its size and position in the record when you select the file formatting parameters using the SE (set parameter) command of the \$IAMUT1 utility. The longer the key, the larger the index. The key should not be longer than necessary but long enough to ensure uniqueness. A shorter key is more efficient than a long key.

ENSURING UNIQUENESS OF THE KEY: To identify each record in an indexed file, each primary key must be unique. If key duplication is possible, the key field must be expanded to ensure that it is unique.

For example, customer name is a key which may involve duplicates. To avoid duplication, lengthen the key field to include other characters such as part of the customer address or the account number. Because the characters in the key must be contiguous, you may need to rearrange the fields in the record.

Another way to eliminate duplication is for you to modify new records dynamically whenever a duplication occurs during loading or processing. One or more characters at the end of the key field can be reserved for a suffix code. Whenever a duplicate occurs, add a value to the suffix and make another attempt to add the record to the file. The result is a file that can contain a sequence of keys such as SMITH, SMITH1, and SMITH2. If you add a suffix, you must use the entire unique key when accessing a record directly.

Providing Access by More Than One Key

To provide good performance with both direct and sequential access, each indexed file is indexed by a single primary key. At times, however, it may be useful to locate records by a secondary key. For example, in a customer file indexed by account number, you might want to locate a record by customer name.

To provide access by a secondary key, you must build a secondary index (a separate file). For a description of setting up secondary indexes, see "Chapter 5. Building a Secondary Index" on page IG-67.


RANDOM AND CLUSTERED RECORD INSERTS

The Indexed Access Method permits records to be added to an existing file. The records are inserted by the Indexed Access Method in the proper


locations according to their key value. This keeps the keys throughout the indexed file in ascending sequence.

Records to be inserted are sometimes required to be distributed throughout the file rather evenly, other times the records to be inserted are in groups.

When there are more individual records to be inserted throughout the file, based on their key value, than there are groups of records to be inserted, this is called random record inserts. The following diagram represents random inserted records among existing records.



Record inserts are considered clustered if most of the inserts occur at only certain places in the file. The following diagram represents clustered inserts by vertically stacked bullets.



DEFINING THE FILE STRUCTURE WITH \$IAMUT1

Defining an indexed file structure is the process of analyzing the file requirements and selecting the appropriate file parameters. This allows you to either precisely define your indexed file or, by proper option selection, \$IAMUT1 will define most of the parameters for you.

\$IAMUT1 is a prompt driven utility. When it is loaded, messages are displayed requesting information to be entered on a keyboard. The responses you enter through the keyboard determine how the utility will operate.

The SE command of the \$IAMUT1 utility permits you to select one of three options for defining your indexed file. The parameter selections are made using the SE command of the \$IAMUT1 utility. The SE (set parameters) command of \$IAMUT1 provides three options for you to choose from to define your indexed file as follows:

1. Option 1 significant parameters - allows you to define an indexed file by supplying a minimum of information. The description of your data records is required and whether you expect random or clustered record insert activity.
2. Option 2 all parameters - allows more flexibility in precisely defining your indexed file but requires more parameters to be supplied.
3. Option 3 parameters from existing indexed data set - can be used when you have an existing indexed file and you wish to use the same parameters for a new indexed file.

\$IAMUT1 Option Selection Guide

Having read the preceding material, you are probably ready to make a choice as to which option you want to use in defining your indexed file. The following table will help you to find the appropriate information, based on your indexed file defining objectives.

Your Objective	Option	Information location
You want the Indexed Access Method to calculate and structure your file	1	See "Designing Indexed Files Using \$IAMUT1 - Option 1" on page IG-21
You want to structure a file and provide specific information for the parameters	2	See "Designing Indexed Files Using \$IAMUT1 - Option 2" on page IG-24
You want the Indexed Access Method to structure a file using the parameters of an existing file	3	See "Designing Indexed Files Using \$IAMUT1 - Option 3" on page IG-57

DESIGNING INDEXED FILES USING \$IAMUT1 - OPTION 1

Option 1 is used if you need to set up your indexed file quickly and easily. You specify only the necessary information and the utility determines the proper values for other parameters. An indexed file generated with this option may not be optimum in terms of storage space performance.

If you want to supply more parameters than are available with this option, or you wish to set up a totally dynamic indexed file, you should see "Designing Indexed Files Using \$IAMUT1 - Option 2" on page IG-24. If you already have an indexed file established and you wish to use those same parameters, you should see "Designing Indexed Files Using \$IAMUT1 - Option 3" on page IG-57.

OPTION 1

The Indexed Access Method utility, \$IAMUT1, option 1 of the SE (set parameters) command, provides you with the opportunity to select only those parameters necessary to set up an indexed file.

The \$IAMUT1 Indexed Access Method utility can be loaded with the Event Driven Executive operator command \$L \$IAMUT1.

When \$IAMUT1 is loaded the first prompt displayed is as follows:

```
[1] ENTER COMMAND (?): SE
```

[1] Entering SE causes the following option list prompt to be displayed.

```
[2] SET FILE DEFINITION PARAMETERS
    0 = EXIT
    1 = SIGNIFICANT PARAMETERS
    2 = ALL PARAMETERS
    3 = PARAMETERS FROM EXISTING INDEXED DATA SET
    ENTER OPTION: 1
```

[2] Respond to this prompt by entering the digit '1'. This response causes a one line prompt from the next prompt sequence to be displayed.

Note: Although the following prompts are displayed one line at a time when using the utility, the entire prompt list is shown for simplicity in describing the parameters.

[3]	SECONDARY INDEX (Y/N)?:	N	
			DEFAULT NEW VALUE
[4]	RECORD SIZE		0:80
[5]	KEY SIZE		0:40
[6]	KEY POSITION		1:1
[7]	BLOCKING FACTOR (RECORDS PER BLOCK)		1:3
[8]	NUMBER OF BASE RECORDS		0:5
[9]	ESTIMATED TOTAL RECORDS		6:20
[10]	TYPE OF INSERT ACTIVITY(C=CLUSTERED,R=RANDOM)		C:R
	DATA SET SIZE IN EDX RECORDS:	12	
	INDEXED ACCESS METHOD RETURN CODE:	-1	
	SYSTEM RETURN CODE:	-1	
[11]	CREATE/DEFINE FILE (Y/N)?:		

[3] The first line asks, are you specifying a secondary index. The response should be N for no, because you are defining the parameters for a primary index file.

[4] The record length shown is 80, however, the entry you will make is the actual record length you want your indexed file records.

[5] Enter the length of your data record field that you are using as the key field. The maximum primary key length is 254.

[6] Enter the position where your primary key field begins. Your data record begins with 1.

[7] Specify the blocking factor (number of records per block) you want your indexed file to have. Remember that when a record is accessed, an entire block is actually read into the system buffer.

[8] Enter the number of base record slots to be defined. This value is the number of records you will load initially. You cannot load more records than this value specifies.

[9] Enter the total number of records you expect this file to contain. records you plan to insert during processing.

[10] Enter the type of record insert activity you expect to have

[11] If you have verified that the parameters you entered are correct, the data set (file) size in EDX records is acceptable, and the return codes are both -1, you can reply Y and you can create and define the file. If you wish to change any of the parameters, reply N and you can reenter the SE command and enter any new values for the parameters.

Replying N terminates the SE function and you can return to this point by reentering the SE command or the DF command (within the same session of \$IAMUT1). The DF command of \$IAMUT1 is described under "DF - Define Indexed File" on page IG-146.

To review the prompts that occur when Y is replied at this point return to the example in "Chapter 2. Using the Indexed Access Method."

DESIGNING INDEXED FILES USING \$IAMUT1 - OPTION 2

Option 2 is used if you have performed an analysis of your file requirements and you want to precisely define your primary indexed file. This option provides a wide range of parameters to allow you to specify your file structure in detail. You can optimize the file structure according to your application requirements for the best storage use and performance.

If you want to supply only the minimum parameters you might want to use option 1 which is described earlier in this chapter under "Designing Indexed Files Using \$IAMUT1 - Option 1" on page IG-21. If you already have an indexed file established and you wish to use those same parameters, you should see "Designing Indexed Files Using \$IAMUT1 - Option 3" on page IG-57.

OPTION 2

The following information is provided so that you can supply the required information to the prompts when defining a primary index file using option 2 of \$IAMUT1. The information is organized in levels of Indexed Access Method blocks. The material should be read sequentially because it provides the information which must be understood in order to apply it to the examples which are placed near the end of this option 2 material.

INDEXED ACCESS METHOD BLOCKS

Indexed files consist of three kinds of blocks:

- Data blocks, which contain records
- Index blocks, which contain pointers to data blocks or lower-level index blocks
- File control blocks, which contain control information

Following is an overview diagram showing the types of blocks and their general relationships to each other in an indexed file.

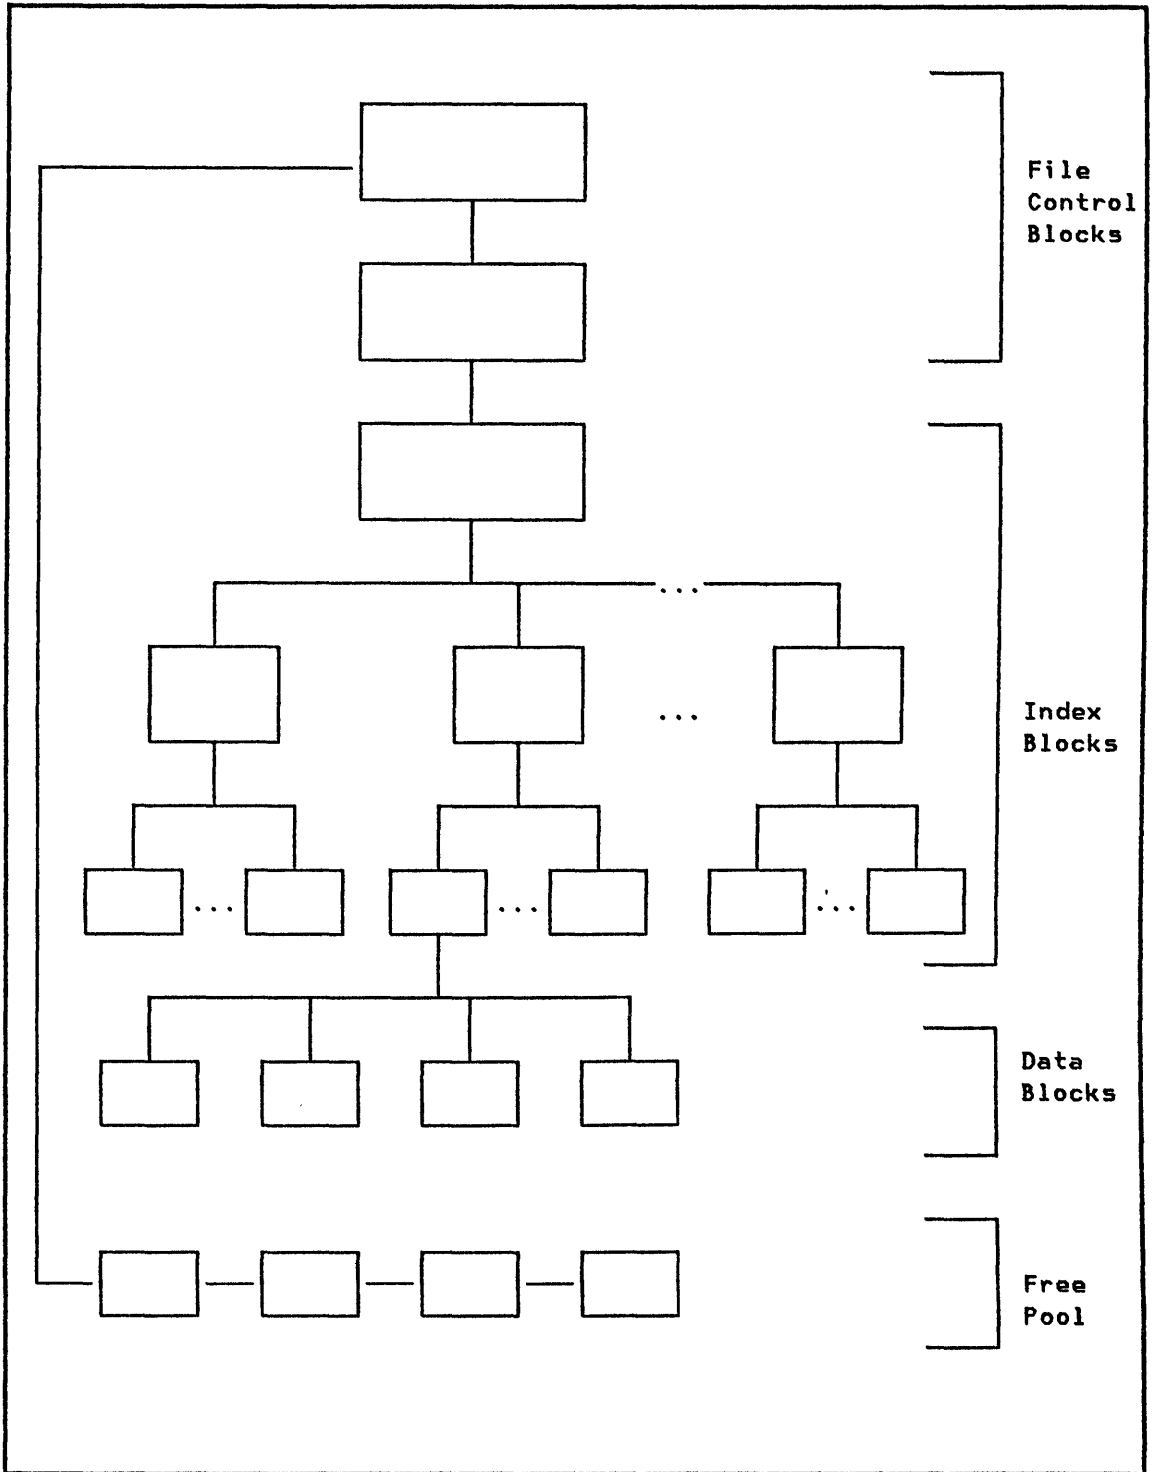


Figure 1. Indexed File Logical Structure

The indexed file is composed of a number of fixed length blocks. The block is the unit of data transferred by the Indexed Access Method between disk/diskette and the central buffer. Block size must be a multiple of

256. A block is addressed by its relative block number (RBN). The first block in the file is located at RBN 0.

Note that the RBN is used only in indexed files by the Indexed Access Method. An Indexed Access Method block differs from an Event Driven Executive record in the following ways:

1. The size of a block is not limited to 256 bytes; its length can be a multiple of 256.
2. The RBN of the first block in an indexed file is 0. The record number of the first Event Driven Executive record in a file is 1.

The size, in 256-byte records, of the file is calculated by the SE command of the \$IAMUT1 utility.

As stated initially, three kinds of blocks exist in an indexed file: data blocks, index blocks, and file control blocks. These blocks are all the same length, as defined by BLKSIZE, but they contain different kinds of information. Data blocks contain data records, index blocks contain index entries, and file control blocks contain control information.

DATA BLOCKS

Each data block contains a header, one or more data records, and it can contain free space for additional data records.

The records in each data block are in ascending order, according to the key field in each record.

Each data block header contains the address of the next sequential data block, providing sequential processing capability.

A data block contains a header followed by data records. The number of records that can be contained in a data block depends on the size of the data block and the size of the record. The header of the block is 16 bytes.

The number of record areas in the block is:

$\frac{\text{block size} - 16}{\text{record size}}$

The result is truncated; any remainder represents the number of unused bytes in the block. For example, if block size is 256 and record size is 80, the data block can accommodate three records and there is no unused area. The key field of the last record slot in an index block is the high key for the data block even if the block is not full. However, if the last record of the block has been deleted, the key field of the last record

slot will contain a key higher than that of any other record in the block. Deletion of a record does not reduce the key range for the block unless the block is emptied. Figure 2, shows the format of a data block.

FREE SPACE

When an indexed file is loaded with base records, free space is reserved for records that may be inserted during processing. There are four kinds of free space: free records, free blocks, reserve blocks, and reserve index entries.

FREE RECORDS: Free records are areas reserved at the end of each data block. The FREEREC parameter of the SE command of \$IAMUT1, specifies the number of free records that are reserved in each data block. The remaining record areas are called allocated records.

For example, if a block contains three data record areas and you specify one free record per block, then there are two allocated records per block. For the layout of a data block containing two allocated records and one free record, see Figure 2.

When records are loaded (file is open in load mode), the allocated records are filled, and the free records are skipped. When additional records are inserted (file is open in process mode), free records are used to hold inserted records.




Figure 2. Data Block Format Example

For an example of specifying FREEREC, see "Example 1: Allocating Free Records" on page IG-43.

FREE BLOCKS: Free blocks follow the allocated data blocks within each cluster. Free blocks have all of their records marked as free records. The FREEBLK parameter of option 2 is used to specify the percentage of blocks that are to be marked as free blocks.

When records are loaded, the allocated record areas in the allocated data blocks are filled, and the free blocks are skipped. During processing, as data blocks become full, a free block provides space for insertions.

For an example of specifying FREEBLK, see "Example 2: Allocating Free Records and Free Blocks" on page IG-45.


SEQUENTIAL CHAINING: Data blocks in an indexed file are chained together by forward pointers located in the headers of data blocks. Only allocated data blocks are included in the sequential chain. Chaining provides for sequential processing of the file with no need to reference the index. When a free block is converted to an allocated block, the free block is included in the chain.

Reserving Space For Record Inserts

If base records are to be loaded and record insertions are expected in random locations throughout the file, use BASEREC to reserve the number of base records. Use some combination of the following parameters: FREEREC to reserve free records in each data block, FREEBLK to reserve free blocks in each cluster (group of blocks), and DYN to provide a free pool.

For example, consider a file with 5 records per block, and 10 data blocks per cluster. Suppose that the file consists of 300 base records and 200 inserts.


If the inserts are distributed evenly throughout the file, the pattern of inserts is:



With this kind of distribution you can specify 2 free records per block to absorb the inserts; no free blocks or free pool are needed.


Of course inserts do not usually occur in such an even pattern. Free blocks help to absorb a concentration of inserts. The more uneven the expected distribution, the greater the free block specification should be.

Suppose the same number of inserts is distributed in this pattern:



With this distribution, specify either 3 free records per block, or 20% free blocks with 2 free records per block.

Now suppose the distribution were more uneven:



In this case a satisfactory mix of free space is 1 free record per block and 40% free blocks. An alternative is to use 1 free record per block and the DYN parameter to hold those record inserts of more than 1 record per block.

Calculating Data Blocks

This calculating information is provided for your convenience if you choose to calculate the number of blocks for a specific file. For reference later there is a summary of all calculations in "Appendix A. Summary of Calculations" on page IG-223. However, \$IAMUT1 automatically calculates the required data blocks based on the parameters you provide. The utility also lists at file definition time (when using the SE command) the number of blocks required according to your parameter values.

The number of allocated data blocks in a file is the specified number of base records (BASEREC) divided by the number of allocated records per data block, with the result rounded up if there is a remainder.

For example, suppose you intend to load 1000 records in an indexed file that is formatted for two allocated records and one free record per block and five allocated blocks and one free block per cluster. The number of allocated blocks in a file is:

number of base records
number of allocated records per block

The number of allocated blocks in this example is $1000/2$ or 500 blocks.

INDEX BLOCKS

An index block contains a header followed by a number of index entries. Each index entry consists of a key and a pointer. The key is the highest key associated with a lower level block; the pointer is the RBN of that block. The number of entries contained in each index block depends on block size and key size. The header of the block is 16 bytes. The RBN field in each entry is 4 bytes. The key field in each entry must be an even number of bytes in length; if the key field is an odd number of bytes in length, the field is padded with one byte to make it even. The number of index entries in an index block is:

block size - 16
4 + key length

The result is truncated; any remainder represents the number of unused bytes in the block.

For example, if block size is 256 and key length is 28, then each index entry is 32 bytes, there are 7 entries in a block, and the last 16 bytes of the block are unused.

PRIMARY INDEX BLOCKS (PIXB)

A set of data blocks is addressed (described) by a single primary index block (PIXB). Each key in the index block is the highest key in the data block that its accompanying relative block number (RBN) addresses. A block is addressed by its RBN. The PIXB and the data blocks it describes are called a cluster.

Clusters

Primary-level index blocks and data blocks are stored together in the file in groups called clusters. Each cluster consists of a primary-level index block and as many data blocks and free blocks as it points to. For example, if there are seven entries in an index block, there are eight blocks in a cluster: one primary-level index block and up to 7 data/free blocks. If reserve blocks have been specified, the blocks represented by

the reserve block entries are not included until insert activity has taken place and the required blocks have been obtained from the free pool. For example, if there are seven entries in an index block and one of the entries is a reserve block entry, the cluster consists of seven blocks (one index block and six data blocks). See Figure 3 for a cluster example.

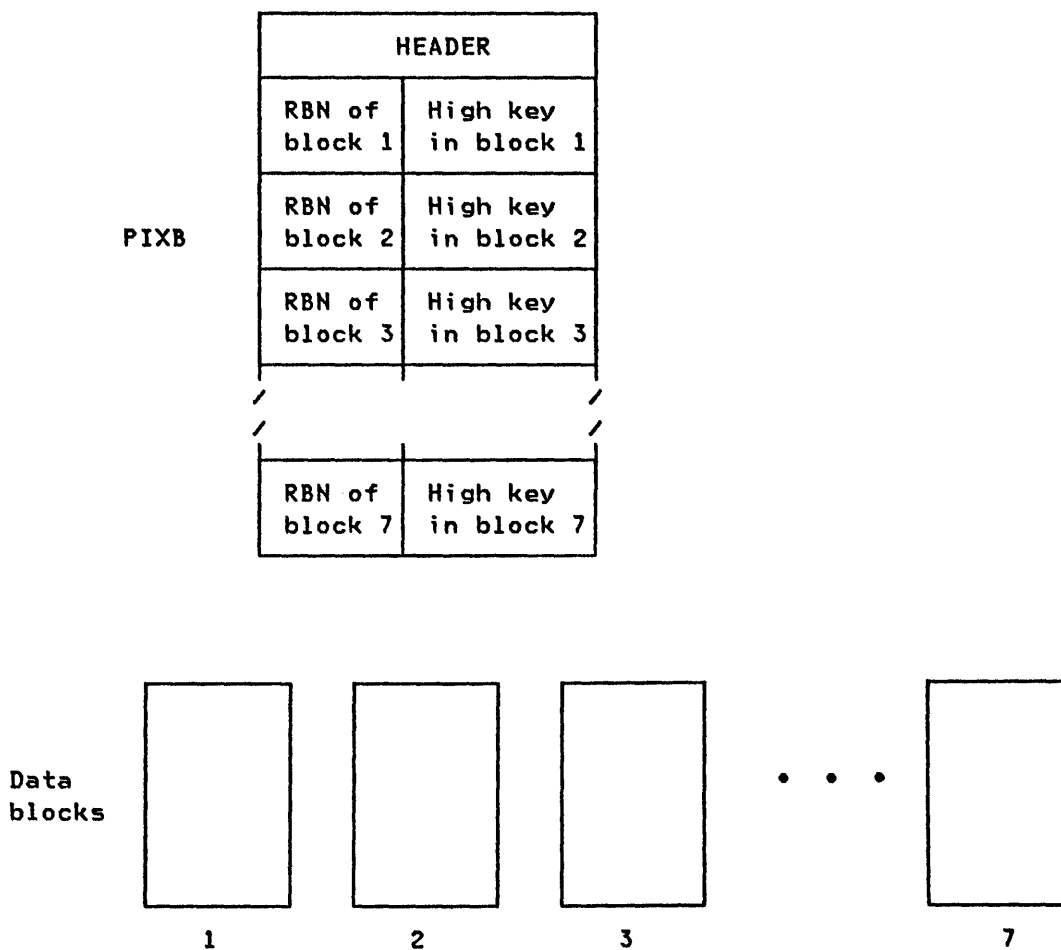


Figure 3. Cluster Example

Primary-Level Index Blocks

Entries in a primary-level index block point to data blocks. Each entry primary-level index block is one of three possible types:

- allocated entry
- free block entry
- reserve block entry

ALLOCATED ENTRY: An allocated entry points to an active data block. The key portion of the entry is initialized to binary ones by the \$IAMUT1 utility. After records have been loaded or written to a data block, the key portion of the entry which points to the data block contains the highest key from the data block.

The pointer portion contains the RBN of the data block. Allocated entries are the first entries in an index block. The number of index entries allocated, when the indexed file is initially created, is the total number of entries per index block, less the number of entries of the other two types (free block entry and reserve block entry).

FREE BLOCK ENTRY: A free block entry points to a free data block. The key portion of the entry contains binary zeros. The pointer portion contains the RBN of the free block. Free block entries follow the allocated entries in the index block. The number of index entries formatted as free entries when the indexed file is initially created is the specified percentage (FREEBLK) of the total number of entries in an indexed block, with the result rounded up if there is a remainder.

RESERVE BLOCK ENTRY: A reserve block entry does not point to a block but is reserved for later use as a pointer to a data block which can be taken from the free pool. Both the key and pointer portions of a reserve block entry are binary zeros. Reserve block entries are at the end of the index block. When a reserve block entry is converted to a used entry, the index block is reformatted to move the entry to the allocated entry area of the block.

Reserve blocks do not exist in the cluster. When all data blocks in a cluster are used and another data block is needed, a data block can be created from the free pool. If the primary-level index block contains a reserve block entry, it is used to point to the record from the free pool. The reserve block entry in the primary-level index block points to the block, and the data block becomes an allocated data block.

The number of index entries initially formatted as reserve block entries is the specified percentage (RSVBLK) of the total number of entries, with the result rounded up if there is a remainder. However, if the number of free block entries plus the number of reserve block entries require all index entries, the number of reserve block entries is reduced by 1, providing at least one allocated entry per index block.

To calculate the number of primary-level index blocks in an indexed file, you must know the initial number of data blocks allocated in the indexed file.

Calculating Clusters

This calculating information is provided for your convenience if you choose to calculate the number of blocks for a specific file. However, \$IAMUT1 automatically calculates the required data blocks based on the parameter values you provide. The utility also lists at file definition

time (when using the SE command) the number of blocks required according to your parameter values.

The number of clusters in a file is the number of allocated data blocks divided by the number of allocated entries in each primary-level index block, with the result rounded up if there is a remainder.

<p><u>allocated blocks</u> allocated entries in each PIXB</p>

Note that in the calculation, if the quotient is not an integer, it is rounded up (rather than truncated) in order to accommodate all of the base records.

The number of free blocks in the file (not including the free pool) is the number of clusters in the file multiplied by the number of free entries in each primary-level index block.

The Last Cluster

The last cluster in the file may be different from the other clusters. It contains the same number of free blocks as the other clusters but only enough allocated blocks to accommodate the records that you have specified with the parameter BASEREC. Because rounding occurs in calculating the number of clusters, a few more allocated records than required may exist in the last allocated block. The last cluster can be a short one because only the required number of blocks are used.

If the number of allocated blocks divided by the number of allocated blocks per cluster leaves a remainder, the remainder represents the number of allocated entries in the primary-level index block in the last cluster. Unused entries in the last primary-level index block are treated as reserve block entries.

The initial number of data blocks is the specified number of base records (BASEREC) divided by the number of allocated records in a data block, with the result rounded up if there is a remainder.

<p><u>BASEREC</u> data records per block</p>
--


The number of primary-level index blocks is the initial number of allocated data blocks divided by the number of allocated entries per primary-level index block, with the result rounded up if there is a remainder.

allocated data blocks allocated entries per primary-level index block

SECOND-LEVEL INDEX BLOCKS (SIXB)

If the file is large enough to require more than one cluster, each PIXB (or cluster) has an entry in a second-level index block (SIXB). The entry in a SIXB contains the address of the PIXB and the highest key in the cluster. The SIXB has the following structure:

HEADER	
RBN of PIXB1	High key in PIXB1
RBN of PIXB2	High key in PIXB2
RBN of PIXB3	High key in PIXB3
RBN of PIXB4	High key in PIXB4



Entries in a second-level index block point to primary-level index blocks. Each entry in a second-level index block is one of two possible types:

- allocated entry
- reserve index entry

ALLOCATED ENTRY: An allocated entry points to an existing primary-level index block. The key portion of the entry is initialized to binary ones by the \$IAMUT1 utility. After records have been loaded or written, the key portion of the entry contains the highest key from the primary-level index block. The pointer portion contains the RBN of the primary-level

index block. Allocated entries are the first entries in the index block. The number of index entries allocated when the indexed file is loaded is calculated as the total number of entries per index block, less the number of reserve index entries.

RESERVE INDEX ENTRY: A reserve index entry does not point to a block but is reserved for later use as a pointer to a primary-level index block that can be taken from the free pool. Both the key and pointer portions of a reserve index entry are binary zeros.

Reserve index entries, in second-level index blocks, provide index space for the index structure to be expanded by adding new primary-level index blocks. These, in turn, can have data blocks associated with them, thus forming new clusters. This process of forming a new cluster is called a cluster split.

For an example of using RSVIX, refer to "Example 4: Allocating Reserved Index Entries" on page IG-51.

Reserve index entries are at the end of the index block. The number of index entries initially formatted as reserve index entries is the specified percentage (RSVIX) of the total number of entries, with the result rounded up if there is a remainder. However, if the number of reserve index entries is the same as the total number of entries in an index block, the number of reserve index entries is reduced by 1, providing at least one allocated entry per second-level index block.

The number of second-level index blocks is the number of primary-level index blocks divided by the number of allocated entries per second-level index block, with the result rounded up if there is a remainder.

$\frac{\text{number of PIXBs}}{\text{allocated entries per SIXB}}$
--

HIGHER-LEVEL INDEX BLOCK (HIXB)

If the file is large enough to require more than one SIXB, the SIXBs in the file are described by one or more higher-level index blocks (HIXB) in the same manner as the SIXB describes PIXBs. There is always one index block that describes the entire file.

The index of an indexed file is constructed in several levels so that, given a key, there is a single path (one index block per level) cascading through the index levels that leads to the data block associated with that key. The index is built from the bottom up. At the lowest level are the primary-level index blocks. At the second level are index blocks containing entries that point to the primary-level index blocks. The highest level of the index structure consists of a single index block.

Entries in a higher-level index block point to index blocks at the next lower level. All entries in higher-level index blocks are allocated entries. The key portion of the entry contains the highest key from the index block of the next lower level. The pointer portion contains the RBN of the next lower level index block. The number of blocks at any higher index level is the number of index blocks at the next lower level divided by the total number of entries per index block, with the result rounded up if there is a remainder.

If the number of index blocks at any level is one, that level is the top level of the index. Although the Indexed Access Method is capable of initially defining and supporting 17 levels of index, an indexed file is formatted with only as many index levels as are required for the number of records. If an indexed file has not been fully loaded and one or more higher index levels have not yet been required, the unnecessary higher levels are not used, even though they exist in the file structure.

INDEX EXAMPLE: Assume that 500 data blocks are allocated to a file and that each primary-level index block contains one free block entry, one reserve block entry, and five allocated entries. Therefore, the total number of primary-level index blocks is 100. Each second-level index block contains one reserve index entry and six allocated entries; therefore, the number of second-level index blocks is 17. The number of entries in higher level index blocks is seven, resulting in three index blocks at the third level and one at the fourth level.

Therefore the file contains a total of 121 index blocks of which 100 are primary-level index blocks, 17 are second-level index blocks, 3 are third-level index blocks, and 1 is a fourth-level index block. This distinction is important because high-level index blocks are located contiguously at the beginning of the file (after the FCB), while primary-level index blocks are scattered throughout the file with the data blocks. Figure 4 on page IG-38 shows the structure of the higher-level index blocks.




Figure 4. High-level Index Structure

FREE POOL

If you specify that you want a free pool, your indexed file contains a pool of free blocks at the end of the indexed file. The file control block contains a pointer to the first block of the free pool, and all blocks in the free pool are chained together by forward pointers.

A block can be taken from the free pool to become either a data block or an index block. The block is taken from the beginning of the chain, and its address (RBN) is placed in the appropriate primary-level index block (if the new block is to become a data block) or in the second level index block (if the new block is to become a primary-level index block), and so on. Any block in the free pool can be used as either a data block or as an index block.

When a data block becomes empty because of record deletions, the data block may return to the free pool (depending on the delete threshold (DELTHR) parameter). If the data block is returned to the free pool, reference to the block is removed from the primary-level index block, and the block is placed at the beginning of the free pool chain.

Calculating the initial size of the free pool consists of the following steps:

- Each reserve block entry in a primary-level index block represents a potential data block from the free pool. The number of data blocks that can be assigned to initial clusters is the number of primary-level index blocks times the number of reserve block entries in each primary-level index block.
- Each reserve index entry in a second-level index block represents a potential primary-level index block from the free pool. The number of primary-level index blocks that can be assigned from the free pool into the index structure set up at file definition time is the number of second-level index blocks multiplied by the number of reserve index entries in each second-level index block.
- Each primary-level index block taken from the free pool consists entirely of empty (reserve block) entries. New data blocks can be taken from the free pool for the entries in the new primary-level index block. The number of data blocks is the number of entries per index block multiplied by the number of new primary-level index blocks (calculated in the previous step).
- The maximum number of blocks that can be taken from the free pool and placed into the index structure set up at file definition time is the sum of the previous three calculations.
- The actual number of blocks in the free pool is determined in one of two ways:
 - The percentage (FPOOL) of the maximum possible free pool as specified by the RSVIX and RSVBLK parameters. The result is rounded up

if there is a remainder. If the DYN parameter is also used, its value is added to the sum.

- The DYN parameter, if specified with no other free space parameters, allocates a free pool of the specified number of blocks.

DELTHR - DELETE THRESHOLD: The percentage (0-99) of blocks to retain in a cluster as records are deleted and blocks made available. This is known as the delete threshold DELTHR. When a block becomes empty, this parameter, if supplied, determines if the block should be returned to the free pool.

FILE CONTROL BLOCK

The file control block (FCB) is the first block in the file (RBN 0); it contains control information.

Indexed files have an FCB Extension as the second block. The FCB Extension contains the parameters used to define the file.

Note: Indexed files built with a version of the Indexed Access Method prior to version 2 do not contain an FCB extension.

You can access the FCB and FCB Extension by either of the following methods:

- using the EXTRACT function in an EDL program
- using the \$VERIFY utility

You can locate the field names in the FCB and FCB Extension by examining a listing of FCBEQU, a copy code module that is supplied as part of the Indexed Access Method. The FCB Extension contains the parameters that were used to set up the file using the \$IAMUT1 SE command. Control information is also contained in block headers; a description of control information is contained in "FCB Extension Report" on page IG-196.

FILE STRUCTURE TYPES

A wide range of file structure is available. You can set up files that vary from the totally dynamic to the highly structured. Whether a file is structured or dynamic depends on the degree to which it uses a free pool.

A free pool is an area in your indexed file which contains a pool of free blocks. The file control block contains a pointer to the first block of the free pool, and all blocks in the free pool are chained together by forward pointers. A block can be taken from the free pool to become either a data block or an index block.

Dynamic files offer the advantage of easy file design and good space utilization. They have the disadvantage of a potential performance decrease.

Structured files offer the advantage of good performance. They have the disadvantage of a more complex file design and greater space requirements.

Either method can result in a need to reorganize the file; the structured approach because the file can run out of space for inserts, and the dynamic approach because of performance considerations.

The type of indexed file to be defined, structured or dynamic, therefore, depends on the file requirements and the efficiency required.

Structured File

A structured file has its base record slots, free space, and the index structure needed to support them built at file definition time by the Indexed Access Method utility using the file structure parameters you specify. The structured file uses little, if any, free pool. The structured file offers better performance than the dynamic file but can result in unused space.

Whether or not a structured file has a free pool depends on whether or not you supply a value for the DYN parameter when the file is defined. When the DYN parameter is used, the FREEREC, FREEBLK, RSVBLK, RSVIX, and FPOOL parameters, if supplied, are also used in establishing the structured free space. The number and types of blocks in a structured file are the result of calculated values you supply as parameters when defining the file. Most of the blocks are not taken dynamically from the free pool as they are needed because they are established at file definition time.

Dynamic File

The higher the degree to which a file uses a free pool, the more dynamic it is; the system builds index and data blocks for you as they are needed.

The Indexed Access Method provides a dynamic file restructuring capability. It makes use of any free pool space the file has, even if the file is mostly structured.

The Indexed Access Method can restructure a file in two ways:

- As records are inserted and additional space is needed in specific areas of the file, blocks are taken from the free pool and become data blocks where needed. If additional index blocks are needed, blocks are taken from the free pool for this purpose as well. Index blocks can be added at any level, and the number of levels of index can increase as needed. This function is performed automatically by the Indexed Access Method on any file that has a free pool associated with it.
- As records are deleted and blocks become empty, they are returned to the free pool. If index blocks become empty (because the blocks under them have been returned to the free pool) they are also returned to

the free pool. This helps to maintain a supply of blocks in the free pool to be used if other areas of the file expand.

For an example of defining a totally dynamic file, see "Example 5 - Defining a Totally Dynamic File" on page IG-54.

USING THE DYN PARAMETER: The DYN parameter can be used to adjust how much the free pool is used. This adjustment varies how dynamic a structured file is.

In a totally dynamic file, the initial file defined consists of only the file control blocks, one primary index block and one data block. The rest of the file is in the free pool.

To define a totally dynamic file, you need to only supply a value for the DYN parameter to allow the rest of the file to be assigned to the free pool.

A dynamic file can be used when the records you want to add to the file are not sorted into ascending key sequence. In that case, you can place the records in the file by inserting them in random sequence. The Indexed Access Method will place them in their proper sequence within the indexed file.

If base records are to be loaded initially and they are sorted in ascending key sequence but insert activity is unknown, you can use a totally dynamic file design. Use the BASEREC parameter to reserve the number of base record slots required. Use the DYN parameter to provide the free pool needed for record inserts.

Note: When a dynamic file has grown to its working size, it should be reorganized for more efficient operation.

OPTION 2 EXAMPLES

The examples which follow are provided to show the option 2 prompts and the effects of certain parameter values. Although the values used are small for simplicity of explanation, they are usually much larger in an actual application. Also a given example does not represent a complete primary index file but addresses a particular part of a file and its associated parameters which we wish to describe at that place in the chapter.

EXAMPLE 1: ALLOCATING FREE RECORDS

The indexed file created using these parameters has only one type of free space, called free records:

```
ENTER COMMAND (?): SE

SET FILE DEFINITION PARAMETERS
0 = EXIT
1 = SIGNIFICANT PARAMETERS
2 = ALL PARAMETERS
3 = PARAMETERS FROM EXISTING INDEXED DATASET
ENTER OPTION: 2
SECONDARY INDEX (Y/N)?: N
PARAMETER  DEFAULT  NEW VALUE
BASEREC      NULL:10
BLKSIZE      0:256
RECSIZE      0:80
KEYSIZE      0:40
KEYPOS       1:1
FREEREC      0:1
FREEBLK      0:0
RSVBLK       NULL:
RSVIX        0:
FPOOL        NULL:
DELTHR       NULL:
DYN          NULL:
TOTAL LOGICAL RECORDS/DATA BLOCK:      3
FULL RECORDS/DATA BLOCK:                2
INITIAL ALLOCATED DATA BLOCKS:         5
INDEX ENTRY SIZE:                       44
TOTAL ENTRIES/INDEX BLOCK:              5
FREE ENTRIES/PIXB:                      0
RESERVE ENTRIES/PIXB(BLOCKS):           0
FULL ENTRIES/PIXB:                      5
RESERVE ENTRIES/SIXB:                   0
FULL ENTRIES/SIXB:                      5
DELETE THRESHOLD ENTRIES:               5
FREE POOL SIZE IN BLOCKS:               0
# OF INDEX BLOCKS AT LEVEL 1:           1

DATA SET SIZE IN EDX RECORDS:           8
INDEXED ACCESS METHOD RETURN CODE:       -1
SYSTEM RETURN CODE:                     -1

CREATE/DEFINE FILE (Y/N)?: N

ENTER COMMAND (?):
```


[1] Because record size was specified as 80 and block size was specified as 256, there are $(256-16)/80 = 3$ records per block. Because FREEREC was specified as 1, there are 2 full (base) records per block and 1 free record per block. Because BASEREC was specified as 10, there are $10/(2 \text{ base records per block})$ or 5 initial allocated data blocks (blocks that contain base records). Because FREEBLK, RSVBLK, RSVIX, FPOOL, and DYN were not specified, there are no free blocks or free pool blocks allocated. One primary index block is needed.

The number of free blocks is calculated as follows: Free entries per PIXB times the number of index blocks at level 1.

The total blocks allocated for this file is:

Initial allocated data blocks	5	
Free blocks	0	
Free pool blocks	0	
Index blocks	1	
File control block	+ 2	
		8 Total

Figure 5 illustrates the format of the indexed file that would result from these SE command parameters.




Figure 5. Indexed File with Free Records

EXAMPLE 2: ALLOCATING FREE RECORDS AND FREE BLOCKS

These parameter specifications will generate an indexed file with two types of free space — free records and free blocks:

```
ENTER COMMAND (?): SE

SET FILE DEFINITION PARAMETERS
0 = EXIT
1 = SIGNIFICANT PARAMETERS
2 = ALL PARAMETERS
3 = PARAMETERS FROM EXISTING INDEXED DATASET
ENTER OPTION: 2
SECONDARY INDEX (Y/N)?: N
PARAMETER  DEFAULT  NEW VALUE
BASEREC      NULL:10
BLKSIZE      0:256
RECSIZE      0:80
KEYSIZE      0:40
KEYPOS       1:1
FREEREC      0:1
[1] FREEBLK    0:10
RSVBLK       NULL:
RSVIX        0:
FPOOL        NULL:
DELTHR       NULL:
DYN          NULL:
TOTAL LOGICAL RECORDS/DATA BLOCK:      3
FULL RECORDS/DATA BLOCK:                2
INITIAL ALLOCATED DATA BLOCKS:         5
INDEX ENTRY SIZE:                       44
TOTAL ENTRIES/INDEX BLOCK:              5
FREE ENTRIES/PIXB:                       1
RESERVE ENTRIES/PIXB(BLOCKS):           0
FULL ENTRIES/PIXB:                       4
RESERVE ENTRIES/SIXB:                   0
FULL ENTRIES/SIXB:                       5
DELETE THRESHOLD ENTRIES:               5
FREE POOL SIZE IN BLOCKS:                0
# OF INDEX BLOCKS AT LEVEL 1:            2
# OF INDEX BLOCKS AT LEVEL 2:            1

DATA SET SIZE IN EDX RECORDS:            12
INDEXED ACCESS METHOD RETURN CODE:        -1
SYSTEM RETURN CODE:                      -1

CREATE/DEFINE FILE (Y/N)?: N
ENTER COMMAND (?):
```

[1] The FREEBLK parameter of 10 causes 10% of the total entries in each index block to point to free blocks. Because KEYSIZE was specified as 40,

the index entry size = 40 + 4 (RBN pointer) and the total entries per index block is $(256-16)/44 = 5$. Thus, 10% of this total rounded up is the number of free entries/PIXB (1). Because there are 5 initial allocated data blocks, one free entry and only 5 total entries per index block, 2 primary index blocks are needed. This causes a second-level index block to be allocated.

The total blocks allocated:

Initial allocated data blocks	5
Free blocks	2
Free pool blocks	0
Index blocks	3
File control block	+ 2
	12 Total

Figure 6 on page IG-47 illustrates the format of the indexed file that would result from these SE command parameters.

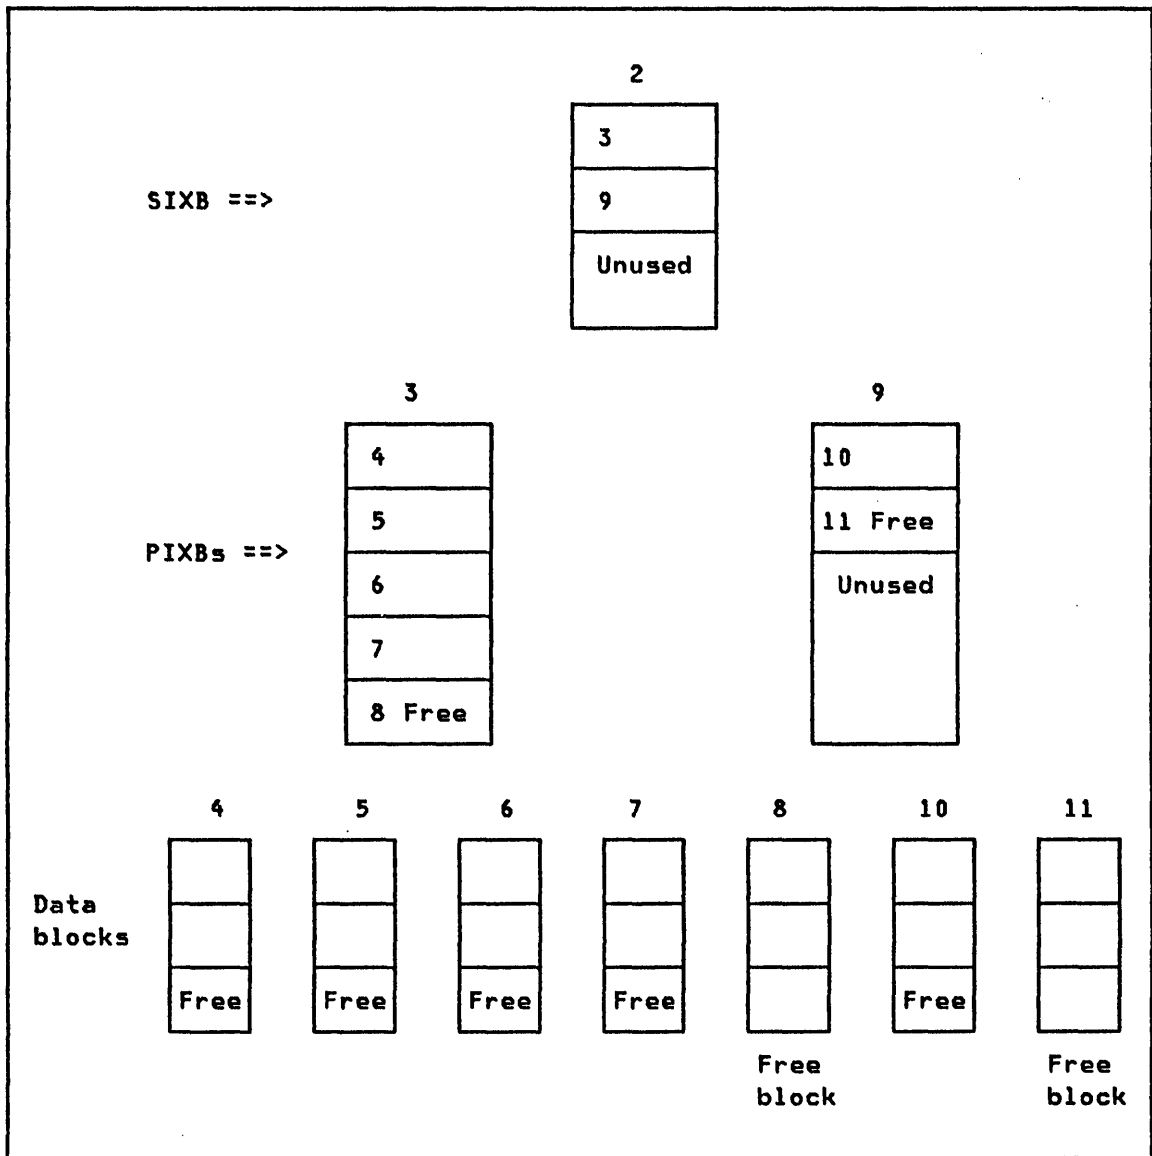


Figure 6. Indexed File with Free Records/Blocks

EXAMPLE 3: ALLOCATING RESERVED DATA BLOCKS

Reserve blocks are allocated using the RSVBLK and FPOOL parameters of the SE command. The following SE command example show the specification of an indexed file with reserved data blocks.

```
ENTER COMMAND (?): SE

SET FILE DEFINITION PARAMETERS
0 = EXIT
1 = SIGNIFICANT PARAMETERS
2 = ALL PARAMETERS
3 = PARAMETERS FROM EXISTING INDEXED DATASET
SECONDARY INDEX (Y/N)?: N
ENTER OPTION: 2
PARAMETER  DEFAULT NEW VALUE
BASEREC      NULL:10
BLKSIZE      0:256
RECSIZE      0:80
KEYSIZE      0:40
KEYPOS       1:1
FREEREC      0:1
FREEBLK      0:10
[1] RSVBLK     NULL:10
RSVIX        0:
[2] FPOOL     NULL:50
DELTHR       NULL:
DYN          NULL:
TOTAL LOGICAL RECORDS/DATA BLOCK:           3
FULL RECORDS/DATA BLOCK:                     2
INITIAL ALLOCATED DATA BLOCKS:              5
INDEX ENTRY SIZE:                           44
TOTAL ENTRIES/INDEX BLOCK:                   5
FREE ENTRIES/PIXB:                           1
RESERVE ENTRIES/PIXB(BLOCKS):                1
FULL ENTRIES/PIXB:                           3
RESERVE ENTRIES/SIXB:                        0
FULL ENTRIES/SIXB:                           5
DELETE THRESHOLD ENTRIES:                    4
FREE POOL SIZE IN BLOCKS:                    1
# OF INDEX BLOCKS AT LEVEL 1:                 2
# OF INDEX BLOCKS AT LEVEL 2:                 1

DATA SET SIZE IN EDX RECORDS:                13
INDEXED ACCESS METHOD RETURN CODE:            -1
SYSTEM RETURN CODE:                          -1

CREATE/DEFINE FILE (Y/N)?: N
ENTER COMMAND (?):
```

[1] In this example RSVBLK was specified as 10. Thus 10% of the total entries in each PIXB will initially be reserved.

[2] Because the total entries per PIXB is 5, 10% of 5 rounded up will cause 1 entry in each PIXB to be reserved. Because there are 2 PIXBs, each with 1 reserve entry, a maximum of 2 free pool blocks can be used. However, since FPOOL was specified as 50%, only half of these blocks (1 block) will be allocated for the free pool.

The total blocks allocated for this file is:

Initial allocated data blocks	5
Free blocks	2
Free pool blocks	1
Index blocks	3
File control block	+ 2
	13 Total

Figure 7 on page IG-50 illustrates the format of the indexed file that would result from these SE command parameters.

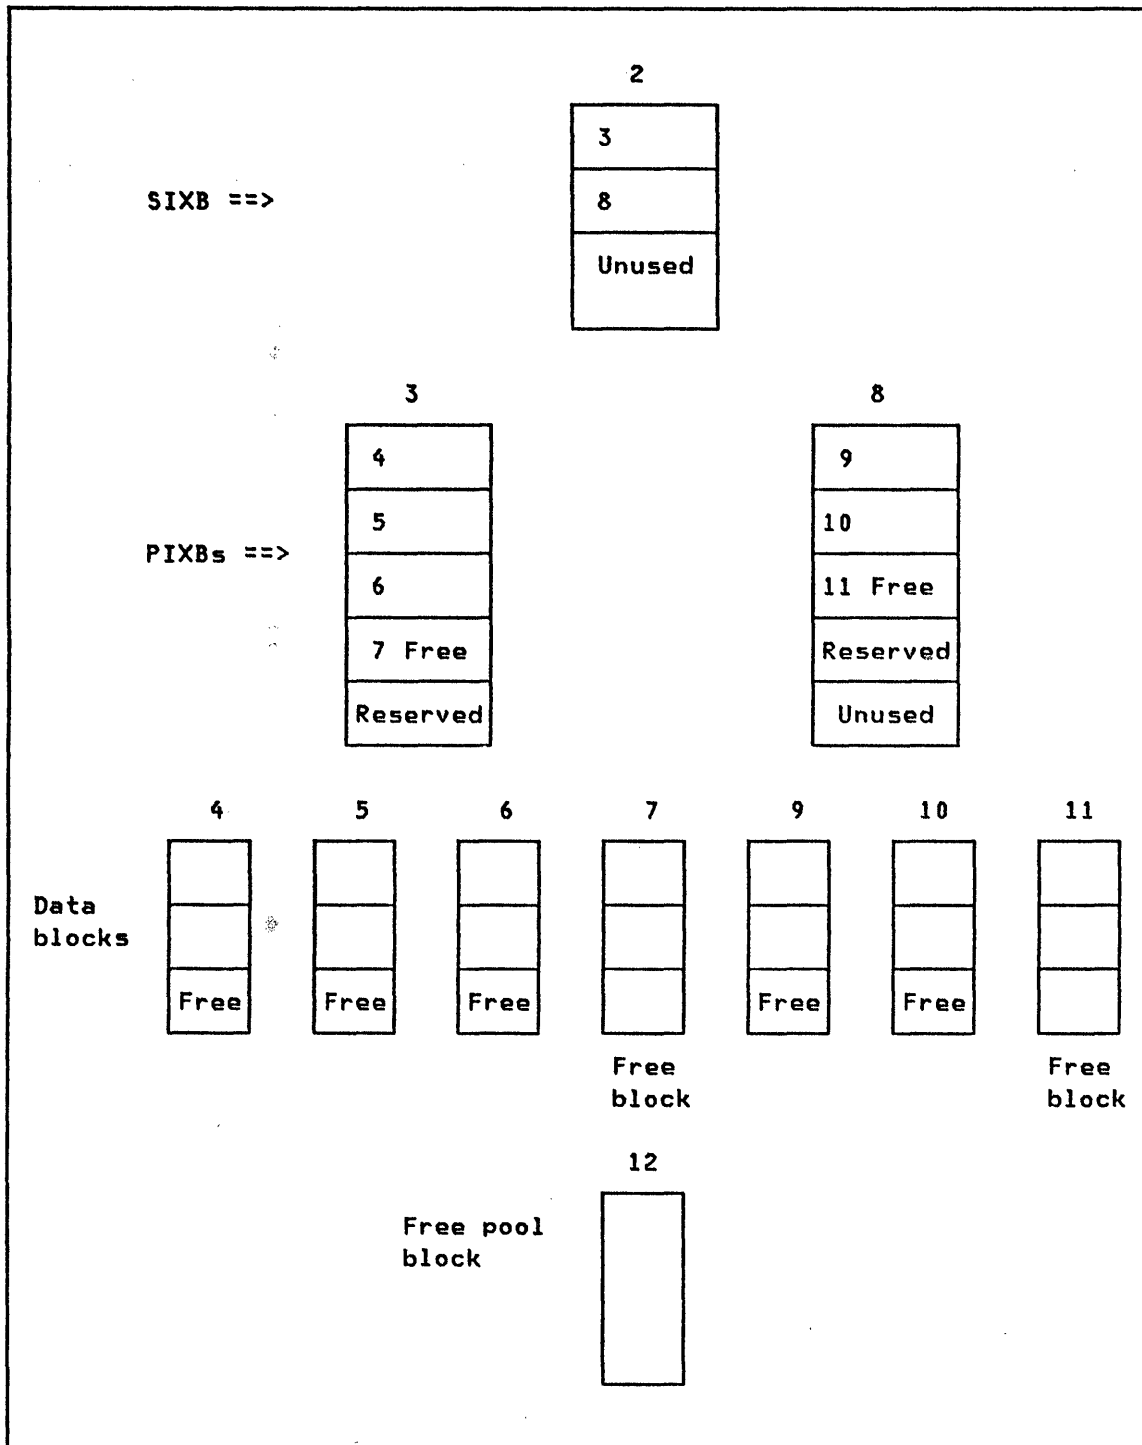


Figure 7. Indexed File with Reserved Data Blocks

EXAMPLE 4: ALLOCATING RESERVED INDEX ENTRIES

In the following example, the index structure is set up to use free pool blocks for index blocks by allocating reserve index entries using the RSVIX parameter.

```
ENTER COMMAND (?): SE

SET FILE DEFINITION PARAMETERS
0 = EXIT
1 = SIGNIFICANT PARAMETERS
2 = ALL PARAMETERS
3 = PARAMETERS FROM EXISTING INDEXED DATASET
ENTER OPTION: 2
SECONDARY INDEX (Y/N)?: N
PARAMETER  DEFAULT  NEW VALUE
BASEREC      NULL:10
BLKSIZE      0:256
RECSIZE      0:80
KEYSIZE      0:40
KEYPOS       1:1
FREEREC      0:1
FREEBLK      0:10
RSVBLK       NULL:10
[1] RSVIX      0:10
[2] FPOOL     NULL:50
    DELTHR    NULL:40
    DYN       NULL:25
TOTAL LOGICAL RECORDS/DATA BLOCK:      3
FULL RECORDS/DATA BLOCK:                2
INITIAL ALLOCATED DATA BLOCKS:         5
INDEX ENTRY SIZE:                       44
TOTAL ENTRIES/INDEX BLOCK:              5
FREE ENTRIES/PIXB:                      1
RESERVE ENTRIES/PIXB(BLOCKS):           1
FULL ENTRIES/PIXB:                      3
RESERVE ENTRIES/SIXB:                   1
FULL ENTRIES/SIXB:                      4
DELETE THRESHOLD ENTRIES:               2
FREE POOL SIZE IN BLOCKS:               29
# OF INDEX BLOCKS AT LEVEL 1:           2
# OF INDEX BLOCKS AT LEVEL 2:           1

DATA SET SIZE IN EDX RECORDS:           41
INDEXED ACCESS METHOD RETURN CODE:       -1
SYSTEM RETURN CODE:                     -1

CREATE/DEFINE FILE (Y/N)?: N
ENTER COMMAND (?):
```


[1] In this example there are still 5 total entries per index block. The 10 RSVIX parameter causes 10% X 5 (rounded up to 1) of the second-level index block (SIXB) entries to be reserved.

In this case, 1 reserve entry is allocated in the SIXB leaving 4 full entries. Because the block pointed to by a SIXB is also an index block (PIXB), blocks in the free pool are allocated for the PIXB and the total number of data blocks it can point to. Thus the total free pool size for these parameters is 1 (reserve entry) + 5 (total entries/PIXB) + 2 (reserve block entries) = 8. Because only 50% of the total possible free pool was requested, 4 of the total free pool blocks plus the 25 blocks specified on the DYN parameter for a total of 29 blocks would be allocated to the free pool.

The total blocks allocated for this file is:

Initial allocated data blocks	5
Free blocks	2
Free pool blocks	29
Index blocks	3
File control block	+ 2
	41 Total

[2] The percentage (0-99) of blocks to retain in the cluster as records are deleted and blocks made available. This is known as the delete threshold DELTHR. When a block becomes empty, it is first determined if the block should be given up to the free pool by checking the response to this prompt. If the block is not given up to the free pool, it is retained in the cluster, either as a free block or as an active empty block. The result of this calculation is rounded up so that any non-zero specification indicates at least one block. The calculation is adjusted to ensure that the cluster always contains at least one block. In this example, the delete threshold was specified as 40%. This results in at least 2 blocks always being retained in each cluster.

If the DELTHR parameter is specified as null (&) and DYN is not specified, DELTHR defaults to the number of allocated blocks in the cluster plus one half of the value calculated by the FREEBLK prompt.

If the DELTHR parameter is specified as null and a value is specified for the DYN parameter, DELTHR defaults to zero.

Figure 8 on page IG-53 illustrates the format of the indexed file that would result from these SE command parameters.

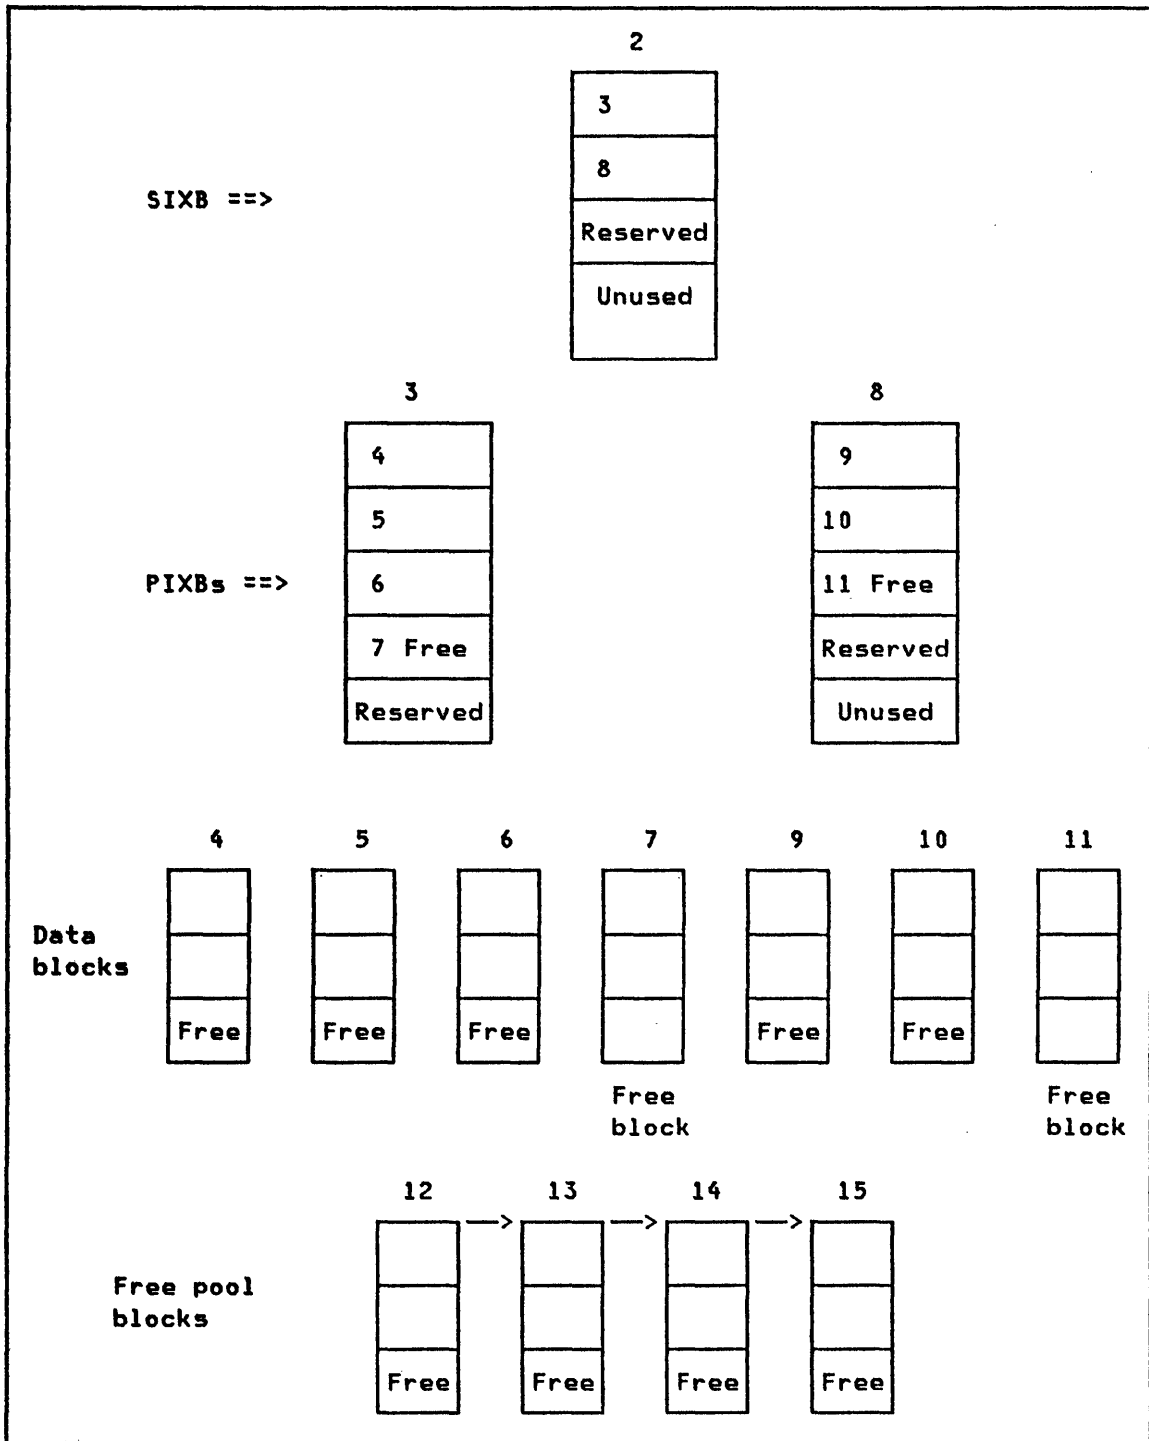


Figure 8. Indexed File with Reserved Index Entries

EXAMPLE 5 - DEFINING A TOTALLY DYNAMIC FILE

To define a totally dynamic file you need only supply the parameters which describe the format of your records within blocks: BLKSIZE, RECSIZE, KEY-SIZE. If the your keys do not begin in position 1 of your records, the KEYPOS parameter must be supplied. The DYN parameter must then be specified in the number of blocks to assign to the free pool.

The following display shows the use of the SE commands of the \$IAMUT1 utility to define a totally dynamic indexed file. Note that the resulting file has only one allocated data block and one index block. The rest of the space is in the free pool as specified by the DYN parameter.

```

ENTER COMMAND (?): SE

SET FILE DEFINITION PARAMETERS
0 = EXIT
1 = SIGNIFICANT PARAMETERS
2 = ALL PARAMETERS
3 = PARAMETERS FROM EXISTING INDEXED DATASET
ENTER OPTION: 2
SECONDARY INDEX (Y/N)?: N
PARAMETER  DEFAULT  NEW VALUE
BASEREC      NULL:
BLKSIZE      0:256
RECSIZE      0:70
KEYSIZE      0:40
KEYPOS       1:
FREEREC      0:
FREEBLK      0:
RSVBLK       NULL:
RSVIX        0:
FPOOL        NULL:
DELTHR       NULL:
DYN          NULL:5300
TOTAL LOGICAL RECORDS/BLOCK:          3
FULL RECORDS/DATA BLOCK:              3
INITIAL ALLOCATED DATA BLOCKS:       1
INDEX ENTRY SIZE:                     14
TOTAL ENTRIES/INDEX BLOCK:            17
FREE ENTRIES/PIXB:                    0
RESERVE ENTRIES/PIXB (BLOCKS):        0
FULL ENTRIES/PIXB:                    17
RESERVE ENTRIES/SIXB:                 0
FULL ENTRIES/SIXB:                    17
DELETE THRESHOLD ENTRIES:             0
FREE POOL SIZE IN BLOCKS:             5300
# OF INDEX BLOCKS AT LEVEL 1:         1

DATA SET SIZE IN EDX RECORDS:         5304
INDEXED ACCESS METHOD RETURN CODE:     -1
SYSTEM RETURN CODE:                   -1

CREATE/DEFINE FILE (Y/N)?: N

ENTER COMMAND (?):

```

The total blocks allocated for this file is:

Initial allocated data blocks	1
Free blocks	0
Free pool blocks	5300
Index blocks	1
File control block	2
	5304 Total

Figure 9 illustrates the format of the indexed file that would result from these SE command parameters.

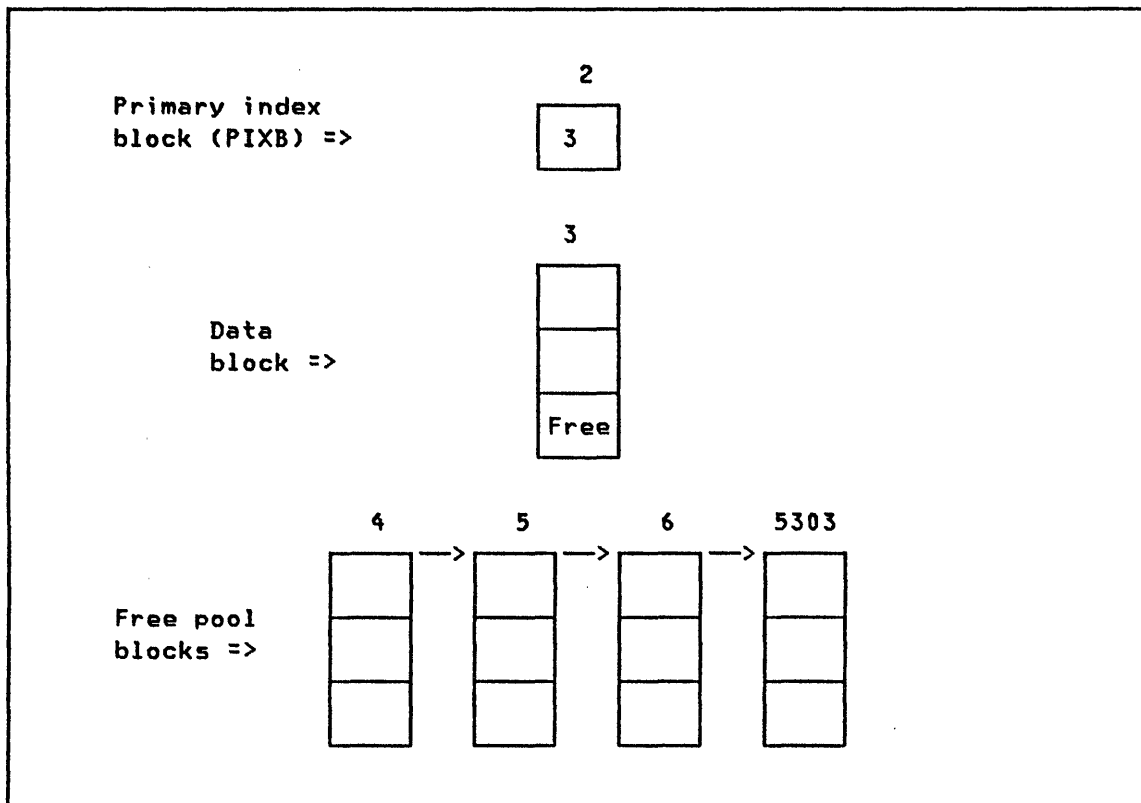


Figure 9. Totally Dynamic Indexed File

DESIGNING INDEXED FILES USING \$IAMUT1 - OPTION 3

Option 3 allows you to define a new file, using the same parameters that were used to create an existing file. Using this option you are not required to manually enter any parameters. You are prompted for the data set name and volume of the existing indexed file followed by the prompt "NEW PARAMATERS EXACTLY SAME AS ORIGINAL PARAMETERS (Y/N) ?". The effects of these two possibilities are described below:

Y The new file to be defined is to appear exactly like the existing file when it was created. In other words, the parameters to be used for defining the new file will be exactly like those of the existing file.

An example of this situation is where you are satisfied with the structure of a currently existing file and now you want to build a similar file and you expect the same type of insert/delete activity.

N The growth of the existing file is to be taken into account in defining the new file. If the total number of records in the existing file do not exceed the number of base records when the file was defined, the existing file parameters will be used without change to define the new file. However, if the number of records in the existing file exceed the number of base records, the parameters for the new file will be adjusted as follows:

- BASEREC will be set as the current number of records in the existing file.
- FPOOL will be set to null.
- DYN will be set to the current number of free pool blocks in the existing file.
- All other parameters will be the same as the corresponding existing file parameters.

Replying N to the prompt "NEW PARAMETERS EXACTLY SAME AS ORIGINAL PARAMETERS (Y/N)?", causes the file size to be adjusted to allow at least as many records to be loaded in the new file as appear in the existing file. This reduces the free pool amount based upon free pool depletion in the existing file.

An example of this situation is where you wish to reorganize a file. The new file should be able to handle as many records as exist in the old file.

Note: The parameters for a primary file must be set from another primary file and parameters for a secondary file must be set from another secondary file.

\$IAMUT1 - OPTION 3

The \$IAMUT1 Indexed Access Method utility can be loaded with the Event Driven Executive operator command \$L \$IAMUT1.

When \$IAMUT1 is loaded the first prompt displayed is as follows:

```
[1] ENTER COMMAND (?): SE
```

[1] Entering SE causes the next prompt to be displayed.

```
[2] SET FILE DEFINITION PARAMETERS
    0 = EXIT
    1 = SIGNIFICANT PARAMETERS
    2 = ALL PARAMETERS
    3 = PARAMETERS FROM EXISTING INDEXED DATA SET
ENTER OPTION: 3
```

[2] Respond to this prompt by entering the digit '3'. This response causes the following prompts to be displayed.

```
[3] SECONDARY INDEX (Y/N)? : N
[3] NAME OF EXISTING INDEXED DATASET (NAME,VOLUME): EMPLFILE,EDX003
[4] NEW PARAMETERS EXACTLY SAME AS ORIGINAL PARAMETERS (Y/N)? Y
    DATA SET SIZE IN EDX RECORDS:           15
    INDEXED ACCESS METHOD RETURN CODE:       -1
    SYSTEM RETURN CODE:                     -1
[5] CREATE/DEFINE FILE (Y/N)? : N
```

[3] Enter the name of the data set and volume whose values you wish this data set to copy.

[4] If all of the parameter values used to define the existing file initially are satisfactory, reply Y. However, if you want to change any of the parameters, based on current file status, or you want to reorganize the existing file, reply N. Replying N will cause the parameter values for BASEREC and FPOOL to be adjusted so that you can load as many records into the new file as are now contained in the existing file.

[5] If you have verified that the parameters you entered are correct, the data set (file) size in EDX records is acceptable, and the return codes are both -1, you can reply Y and the file can be defined and created. If

you wish to change any of the parameters, reply N and you can reenter the SE command and enter any new values for the parameters.

DEFINING, CREATING, AND LOADING A FILE - SUMMARY

This chapter has presented the structure, content and principles of primary index files. Several examples have been used to show what results given parameter values when defining a primary index file.

In those examples the use of the SE command of \$IAMUT1 was used extensively. In replying to the SE prompt "DEFINE/CREATE FILE (Y/N)?:", N was used in this chapter. This allows you to reenter the SE command and go through the prompt sequence again, changing any parameter values as required.

To see the result of replying Y, see the example used in "Chapter 2. Using the Indexed Access Method" on page IG-5. When you reply Y to the DEFINE/CREATE prompt, you enter the function called defining the file. You can enter the define file directly anytime that \$IAMUT1 is loaded by responding with DF to the prompt "ENTER COMMAND (?):." Using the DF (define file) command is described in detail under "DF - Define Indexed File" on page IG-146.

When you reply Y to the prompt, "INVOKE LOAD(L), REORGANIZE(R) OR END(E) AFTER CURRENT FUNCTION ?:", you are given the opportunity to enter the \$IAMUT1 functions of load, reorganize, or end. While in the SE function, load, reorganize, and end can be entered by replying with the letters L, R, or E, respectively. However, these functions can be entered directly from the prompt "ENTER COMMAND (?):" with LO for load, RO for reorganize, or EN for end.

Using the LO (load) command is described in detail under "LO - Load Indexed File" on page IG-162.

Using the RO (reorganize) command is described in detail under "RO - Reorganize Indexed File" on page IG-171.

Entering EN (end) terminates the current session of the SE command of \$IAMUT1. Entering EN to the prompt "ENTER COMMAND (?):" will then terminate the \$IAMUT1 utility.

CHAPTER 4. LOADING THE PRIMARY INDEX FILE

This section describes the process and methods of loading a file.

You can use two methods to load base records:

- The \$IAMUT1 utility
- An application program

The methods are described in the following sections.

LOADING THE PRIMARY INDEX FILE

The Indexed Access Method uses two modes to place records into an indexed file:

1. **Load mode:** records are loaded sequentially in ascending order by key, skipping any free space. The records loaded are called **base records**. Each record loaded must have a key higher than any key already in the file.
2. **Process mode:** records are inserted in their proper key position relative to records already in the file. Records are inserted using the free space that was skipped during loading or, if a record has a new high key, it is placed in a base record after the last loaded record. If no base records are available, it is placed in the free space after the last loaded record.

The total number of base records that can be loaded is established when the indexed file is defined by the \$IAMUT1 utility. It is not necessary, however, to load all (or any) base records before processing can begin. The file can be opened for loading some of the base records, closed and then reopened for processing (including inserts), and later opened for loading more base records. Figure 10 on page IG-62 illustrates this sequence.

Note: Programs written in COBOL are an exception to this: COBOL programs can use load mode only once for any given indexed file. Therefore, all base records loaded in load mode must be loaded together. Base records loaded later must be inserted in process mode (with slower performance).

The amount of free space for inserts (if any) is specified using the \$IAMUT1 utility when the indexed file is built. This free space can be distributed throughout the file in the form of free records within each data block, free blocks within each cluster, and in a free pool at the end of the file.

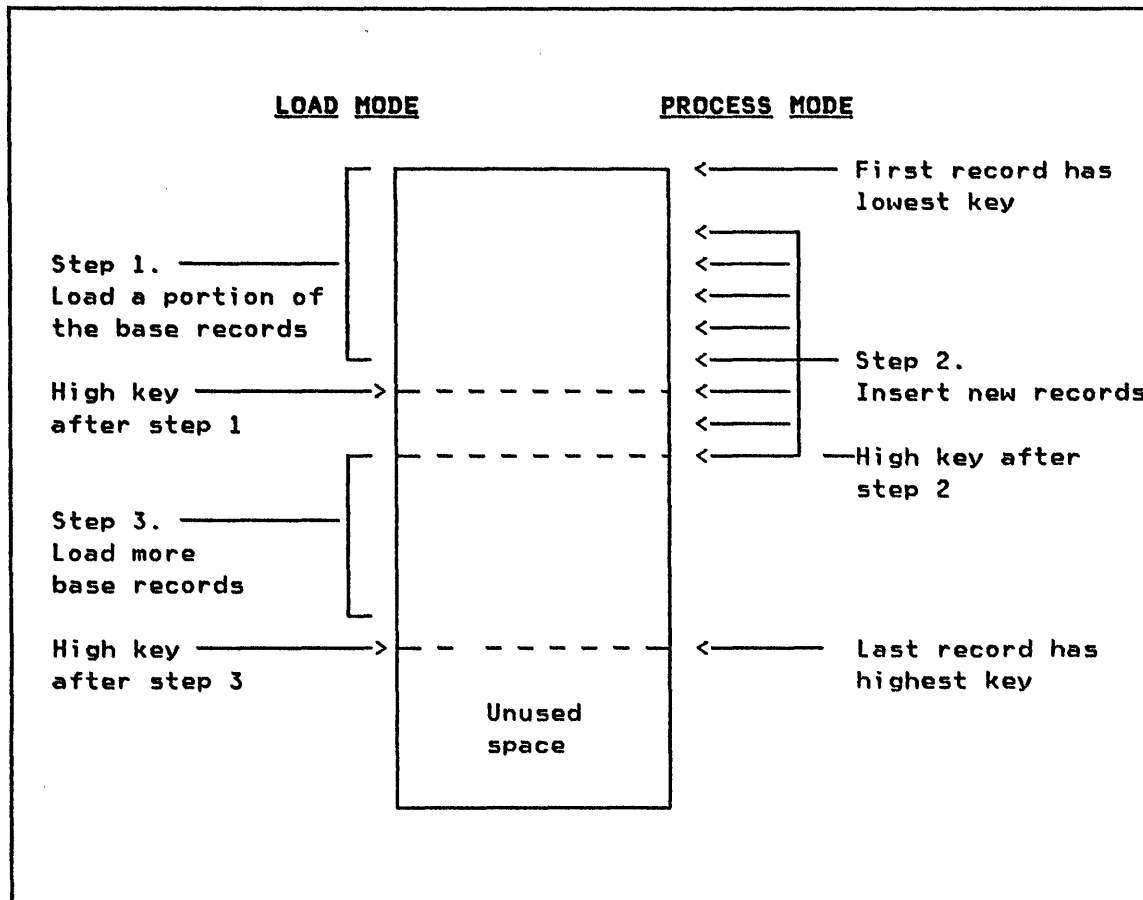


Figure 10. Loading and Inserting Records

LOADING BASE RECORDS USING \$IAMUT1

After the indexed file has been defined by the \$IAMUT1 DF command, you can load base records from a sequential file into the indexed file. Loading the file can be done directly by responding Y to the prompt "INVOKE LOAD(L), REORGANIZE(R) OR END(E) AFTER CURRENT FUNCTION?", when defining the file, or by using the LO command after the file has been defined. The data in the sequential file must be in ascending order by key.

To load base records using \$IAMUT1, do the following:

1. Prepare a sequential file for input to the indexed file
2. Load the sequential file into the indexed file

Preparing Input for the Indexed File

Select one of the following methods to prepare the input in a sequential file to be loaded into an indexed file:

- If your data records are 72 bytes or less, use one of the text editors to enter your data or one of the communications utilities to get the data into an Event Driven Executive sequential file. In either case, you must know the record format used by the utility. The utilities put two 80-byte records in each 256-byte Event Driven Executive record. The first record begins at location 1, and the second record begins at location 129. This results in a blocked sequential file which can be used to load the indexed file when using the LO command of \$IAMUT1. (A detailed description of the LO command is under "LO - Load Indexed File" on page IG-162.) Specify 128 for the input record length and 256 for the input block size.
- If your records have more than 72 bytes of data, you must create a program that accepts the data records and writes them to a disk, diskette, or magnetic tape file.

The data must be in ascending order, based upon the field you use as the key.

Loading an Indexed File from a Sequential File

The procedure for loading an indexed file from a sequential file is:

1. Invoke \$IAMUT1 using the system command \$L.
2. If you want a hard copy of the terminal prompts and responses, enter an EC command. Respond to the prompt with a Y. This will print all further prompts and responses of \$IAMUT1 on the \$SYSPRTR device and your terminal. If a hard copy is not required, omit this step.
3. Enter the LO command.

Respond to the following prompts with your data set information.

```
ENTER COMMAND (?): LO
LOAD ACTIVE
ENTER OUTPUT DATASET (NAME,VOLUME):
$FSEDIT FILE RECSIZE = 128
INPUT RECORD ASSUMED TO BE      80 BYTES. OK?:
ENTER INPUT BLOCKSIZE (NULL = UNBLOCKED):
ENTER INPUT DATASET (NAME,VOLUME):
LOAD IN PROCESS

END OF INPUT DATASET
ANY MORE DATA TO BE LOADED?: N
      6 RECORDS LOADED
LOAD SUCCESSFUL
```

4. Enter the EN command to end \$IAMUT1. Your program is now loaded and you can process the data with your application program.

LOADING BASE RECORDS FROM AN APPLICATION PROGRAM

Base records are records placed into an indexed file in ascending new high key sequence. That is, if a record added to the file has a key higher than any other record in the file, it is placed in a base record slot. Base records are placed in the base record slots reserved for them by use of the BASEREC parameter. You can use either the \$IAMUT1 LO command or an application program to load the base records.

Base records must be loaded in ascending order by key. If you are writing your own program to load the file, use a LOAD request to connect the file to load base records. Then issue a PUT for each record. When the desired records have been loaded, issue a DISCONN request to terminate the load procedure. The only requests that can follow a LOAD request are: PUT, EXTRACT, and DISCONN.

You can also insert base records in process mode by using a PROCESS request to connect the file, followed by a PUT request for each record to be loaded. Loading records in load mode with an application program is discouraged because of slower performance.

Unless the base record loading program is written in COBOL, it need not load all base records at one time. A file that already contains records can be reconnected to load more records, but the key of each new record must be higher than any key already in the file.

COBOL programs must either load all the base records in load mode at once (because only one use of load mode is allowed on a given file) or insert the records in process mode as needed.

The limit on base records as specified on the SE command of the Indexed Access Method utility program (\$IAMUT1) cannot be exceeded. If you attempt to load a record after the last allocated record area has been filled, an end-of-file condition occurs.

LOADING BASE RECORDS FROM A SEQUENTIAL FILE IN RANDOM ORDER

In order to load base records from a sequential file where keys are in random order, code an EDL program to open the indexed file in load mode. Load the SORT/MERGE program with an output exit routine specified. Write (PUT) each record to the indexed file as it is received in the output exit routine from SORT/MERGE. The output exit routine can also screen out duplicates or other unwanted records. For information on using the SORT/MERGE Program Product, refer to Sort/Merge Programmer's Guide.

CHAPTER 5. BUILDING A SECONDARY INDEX

Indexed files, like most data record files, can be a common base for many applications. You can assign **secondary keys** in your indexed files for greater flexibility in accessing records in indexed files.

Secondary keys are accessed through a **secondary index** (a separate file). Your application program requests records by their secondary key and secondary index file name. The secondary index is used to retrieve the record by its secondary key from the primary index file.

You can have more than one secondary index for a given primary index file. In order for the Indexed Access Method to know the relationships between secondary indexes and primary index files, you must create and maintain a **directory** with that information.

SECONDARY KEYS

Secondary keys are not required to be unique; different records in an indexed file can have the same key values in their secondary key field.

The secondary key can be any field within your data record that you select, however, it must meet the following requirements:

- the selected field must start at the same location in each record
- all portions of the key field must be contiguous
- the secondary key length cannot exceed 250 bytes

In a secondary index, the Indexed Access Method assigns a **sequence number** to each secondary key. The sequence number shows the sequence of loading or inserting secondary index entries.

A sample layout of a secondary index record follows:

Secondary Key	Sequence Number	Primary Key	Relative Block Number (primary)
SMITH	0001	12345AB	RBN

THE DIRECTORY

In order for the Indexed Access Method to know the relationships between secondary indexes and primary index files, you must create and maintain a

directory with that information. The directory describes all indexed files in the system which are either secondary indexes, or primaries which have secondary indexes associated with them. Primary index files which do not have secondary indexes associated with them are not in the directory. Use the \$IAMUT1 utility to create and maintain the directory.

The directory name is \$IAMDIR and it resides on the IPL volume.

The directory contains one or more groups of entries. Each group begins with an entry for the primary file and is followed by an entry for each secondary index which references that primary file.

You have the responsibility of maintaining the directory using the \$IAMUT1 utility.

Each entry in the directory contains the following information:

- file name
- volume name
- primary file or secondary index indicator
- independent indicator
- invalid index indicator (secondary entry only)
- automatic update indicator (secondary entry only)

FILE NAME: The file name is the data set name supplied when the primary index file or secondary index entry is inserted in the directory.

VOLUME NAME: The file location is the volume label name where the index resides that this entry is for.

INDEPENDENT PROCESSING INDICATOR: Each entry in the directory contains an independent indicator. Independent means that the file is to be treated as an independent file without regard to associated primary or secondary files. If the independent indicator is set on for a file that is explicitly opened, the automatic update indicator is ignored.

In the case of a secondary index, this means that records retrieved are internal secondary index records, not data records from the primary file. In addition, independent means that any modification to the file (either primary or secondary) will not be reflected in its associated files. Also any changes made in a secondary index will not be reflected in the associated primary or other secondary index files.

In the case of a primary entry, any modification to the primary file will not be reflected in the associated secondary index files.

INVALID INDICATOR: The invalid indicator is initially turned on in the directory, by the directory function of \$IAMUT1, when the secondary entry is inserted in the directory.

A secondary index entry is marked invalid until the secondary index has been loaded.

The load function of the utility turns off the invalid indicator.

If you build the secondary index with an application program, you must also turn off the indicator. The UE subcommand of the DR function in \$IAMUT1 is used to turn off the invalid indicator, after you have successfully loaded your secondary index.

AUTOMATIC UPDATE INDICATOR: Each secondary index entry in the directory contains an automatic update indicator. Any modification to the primary file (either directly or through any secondary index activity) results in an automatic update to all secondary indexes whose automatic update indicator in the directory was specified with Y. Thus, a secondary index flagged as auto-update can be thought of as "dynamic." Each secondary index remains open until all users of it have closed. However, if the independent indicator is set on for a file that is explicitly opened, the automatic update indicator is ignored.

If the automatic update indicator was specified as N, changes are not reflected in that secondary index. This would be a "static" index. The assumption is that a static index would be rebuilt periodically to bring it up to date.

ALLOCATING AND INSERTING ENTRIES IN A DIRECTORY

Although the Indexed Access Method references the directory, it never modifies the directory. The one function that is performed on the directory automatically is that the secondary load option sets the invalid indicator off following successful completion.

To define the existence of a secondary index, use \$IAMUT1 to perform the following two steps:

1. allocate a directory using the DR (directory function) of \$IAMUT1
2. establish the fact that a secondary index will exist by making an entry in the directory using the IE (insert entry) command of \$IAMUT1

Remember that primary index file entries precede their associated secondary index entries. The \$IAMUT1 Indexed Access Method utility can be loaded with the Event Driven Executive operator command \$L \$IAMUT1.

When \$IAMUT1 is loaded the first prompt displayed is as follows:

```
[1] ENTER COMMAND (?): DR
```

[1] Entering DR causes the following prompt sequence.

```

[2] ENTER DIRECTORY COMMAND (?): AL
[3] MAX # OF DIRECTORY ENTRIES: 10
[4] THE DIRECTORY DS REQUIRES 1 EDX RECORDS, CONTINUE (Y|N|EN)? Y
[5] DIRECTORY DATA SET ALLOCATED: $IAMDIR,EDX002

```

[2] Responding to this prompt with AL (allocate) causes a directory allocation sequence to begin.

Note: The allocation sequence is only required the first time you set up secondary indexes. Future entries can be added using (IE) insert directory function.

[3] Reply with maximum number of directory entries you want allocated for the directory. You will need one entry for each secondary index and one entry for each primary that has a secondary index associated with it. A null response will allocate the maximum (default) of 47 entries.

[4] Based on your previous response, the size of the required directory is calculated and you are informed of the number of EDX records required to allocate your requested directory. You are also given three options as follows:

1. Y - the opportunity to continue the directory allocation
2. N - do not allocate the directory; allow me to change the size of the directory
3. EN - end the allocate function; return to [1] of the DR function of \$IAMUT1 to enter another command.

[5] Because [Y] was replied, the directory is allocated. If the directory is allocated successfully, you are informed that it has been allocated, the name of the directory of course is \$IAMDIR, and the IPL volume where it is always allocated is displayed.

Note: The allocation sequence is only required the first time you set up secondary indexes.

The prompt sequence continues.

```
[6]  ENTER DIRECTORY COMMAND (?):  IE
[7]  (DSNAME,VOLUME):  EMP#,EDX002
[8]  IS THIS A SECONDARY ENTRY (Y/N) N
[9]  DIRECTORY INSERT SUCCESSFUL
[10] ENTER DIRECTORY COMMAND (?):  IE
```

[6] Replying **IE** (insert entry) allows you to insert entries into the directory. A primary entry must be inserted before its associated secondary index entries.

Note: Primary files may exist at this time, however, secondary indexes cannot be created until an entry for it has been inserted in the directory.

[7] Your data set name and volume name where your primary index file or secondary index resides. The volume name is not required if the data set is on system volume such as EDX002.

[8] This prompt lets \$IAMUT1 know whether to set the primary or secondary entry indicator. Reply Y for a secondary index entry, or N for a primary index entry.

[9] This message informs you that the entry has been successfully inserted into the directory.

[10] At this point you can end the directory function by responding to the prompt with EN, or reply any other directory function.

Because this was a primary entry we can now respond with IE and insert secondary directory entries. In this case, secondary entries are being made and that is why we responded with IE and caused the prompts to continue as follows:

```
[11] (DSNAME,VOLUME): NAME,EDX002
[12] IS THIS A SECONDARY ENTRY (Y/N)? Y

[13] ASSOCIATED PRIMARY ENTRY (DSNAME,VOLUME): EMP#,EDX002
[14] AUTO-UPDATE (Y/N)? Y

[15] DIRECTORY INSERT SUCCESSFUL

[16] ENTER DIRECTORY COMMAND (?): IE
```

[11] The secondary index data set name is NAME on volume EDX002 and therefore, the volume name could have been omitted.

[12] Because this is an entry for a secondary index, the correct reply is Y. At this point in the prompt sequence the prompts change from the previous sequence because of the positive reply, Y.

[13] The associated primary entry data set name, which the previous entry sequence (5 - 9) was for, is EMP#,EDX002. This is the point where the secondary indexes establish their association to the primary index files for which the secondary index is built.

[14] The response to this prompt establishes whether automatic update option is to be effective for this secondary index. For a description of automatic update, see "Automatic Update Indicator" on page IG-69. The recommended response is Y, also if a null entry is supplied, the default is Y (yes).

[15] You are informed when the insert is successfully completed.

[16] As seen previously, you again have the option of selecting another directory function. In this description, IE was again selected to insert the following two secondary index entries.

A second secondary index entry named CITY,EDX002, is inserted for the associated primary index file named EMP#.

```
(DSNAME,VOLUME): CITY,EDX002
IS THIS A SECONDARY ENTRY (Y/N) Y
```

```
ASSOCIATED PRIMARY ENTRY (DSNAME,VOLUME): EMP#,EDX002
AUTO-UPDATE (Y/N)? Y
```

```
DIRECTORY INSERT SUCCESSFUL
```

```
ENTER DIRECTORY COMMAND (?): IE
```

A third secondary index entry named LEVEL,EDX002, is inserted for the associated primary index file named EMP#.

```
(DSNAME,VOLUME): LEVEL,EDX002
IS THIS A SECONDARY ENTRY (Y/N) Y
```

```
ASSOCIATED PRIMARY ENTRY (DSNAME,VOLUME): EMP#,EDX002
AUTO-UPDATE (Y/N)? N
```

```
DIRECTORY INSERT SUCCESSFUL
```

The following example uses a different directory function: LE (list directory entries). This example shows the directory which was just allocated and four entries inserted; one primary and three secondaries.

```
[1] ENTER DIRECTORY COMMAND (?): LE
```

```
[2] ENTRY (DSNAME,VOLUME) BLANK=ALL:
```

```
[3]          PRIMARY INDE-          AUTO
DSNAME  VOLUME DATA SET PENDENT  INVALID  UPDATE
[4] EMP#    EDX002 YES      NO      ****   ****
[5] NAME    EDX002 NO       NO       YES    YES
[6] CITY    EDX002 NO       NO       YES    YES
[7] LEVEL   EDX002 NO       NO       YES    NO
```

```
[8] NUMBER OF DIRECTORY ENTRIES USED =      4
```

```
[9] NUMBER OF AVAILABLE ENTRY SLOTS =      6
```

```
DIRECTORY LIST COMPLETED
```

[1] The DR (directory) subcommand LE prints specified directory statistics.

[2] Respond to this prompt with the specific data set name and volume you wish the statistics listed for, or press the Enter key with no DSNAMES or VOLUME name specified to list the entire directory. This request is for all entries in the directory which was just allocated and inserts made in the previous examples.

[3] Column headings for the listed information from the directory showing the following information:

- data set name that the statistics are for
- volume name where the data set resides
- whether this is a primary or secondary index
- is the independent indicator on for the named data set (yes or no)
- is the invalid index indicator on for the named data set (yes or no)
- is the auto-update indicator on for the named data set (yes or no)

[4] For the primary index file (data set) named EMP#, on volume EDX002, the independent indicator is off, there is no invalid indicator for a primary file, there is no auto-update indicator for a primary file. Modifications are always made to the primary index file if the independent indicator is not on.

[5] For the secondary index named NAME, on volume EDX002, the independent indicator is off, the invalid index indicator is on because the index has not been loaded, the auto-update indicator is on as requested when the entry for this secondary index was inserted.

[6] Same statistics as previous data set.

[7] For the secondary index named LEVEL, on volume EDX002, the independent indicator is off, the invalid indicator is on (index is not loaded), and the auto-update indicator is off.

[8] There were 10 entries allocated and 4 inserts (one primary and three secondaries).

[9] The resulting empty directory slots for additional inserts is six.

SECONDARY INDEX

Depending upon your need, you may have one or several secondary indexes for a given primary index file. A secondary index is built for a specific primary index file and cannot be used with any other file. Each secondary index is a separate Indexed Access Method file.

Application programs accessing indexed records by their secondary key are required to open the secondary index and access the records using the secondary key. When primary index records are updated, inserted or deleted, some or all secondary indexes associated with that primary index file can be updated automatically by the Indexed Access Method, according to the options selected when the secondary index directory is set up.

Setting up a Secondary Index

To provide access by a secondary key, you must build a secondary index. The secondary index must have a unique file name.

To set up a secondary index, you must do the following using \$IAMUT1:

1. create the secondary index
2. load the secondary index

DEFINING AND LOADING A SECONDARY INDEX

Your secondary index should be structured so that the base records parameter is equal to or greater than the number of records in the primary index file. This will assure that when you build your secondary index, it will be large enough to hold at least as many records as there are in the primary index file.

Note: If the associated primary index file, for which the secondary is being defined, is an existing Version 1 created file, you must use \$VERIFY to update the record counts before defining the secondary file.

The key size and key position specified for the secondary index must be the key size and starting position of the secondary key within the primary index record.

You can create a secondary index the same way you create a primary index file, using the \$IAMUT1 utility SE (and DF) commands. The utility prompts you requesting whether the secondary index being defined is also to be loaded. If YES is specified, the utility does the following:

1. creates the secondary index but does not format it
2. opens the primary file, reads the records sequentially, and extracts the primary and secondary keys from each record, retaining the relative data block address (RBN) of each record
3. invokes the Sort/Merge Program Product to sort by secondary key (and by primary key within secondary)
4. opens the secondary index, formats the sorted keys, their sequence numbers which are now assigned, and the relative data block addresses of the primary file data records into blocks

5. writes the blocks into the secondary index

Before a secondary index can be loaded, it must have been defined using \$IAMUT1. A secondary index can be deleted, then created and loaded again at any time. If a primary file has more than one secondary index, each must be created and loaded separately.

\$IAMUT1 Option Selection Guide

Having read the preceding material, you are probably ready to make a choice as to which option you want to use in defining your secondary index. The following table will help you to find the appropriate information, based on your secondary index defining objectives.

Your Objective	Option	Information location
You want the Indexed Access Method to calculate and structure your index	Option 1	See "Option 1" on page IG-77
You want to structure your secondary index using specific parameters	Option 2	See "Option 2" on page IG-79
You want to structure your secondary index using the parameters of an existing secondary index	Option 3	See "Option 3" on page IG-82

EXAMPLE 1: DEFINING A SECONDARY INDEX USING \$IAMUT1

The Indexed Access Method utility, \$IAMUT1, option 1, provides you with the opportunity to select only those parameters necessary to set up a secondary index.

OPTION 1

The \$IAMUT1 Indexed Access Method utility can be loaded with the Event Driven Executive operator command \$L \$IAMUT1.

When \$IAMUT1 is loaded the first prompt displayed is as follows:

```
[1] ENTER COMMAND (?): SE
```

[1] Entering SE causes the following option list prompt to be displayed.

```
[2] SET FILE DEFINITION PARAMETERS
    0 = EXIT
    1 = SIGNIFICANT PARAMETERS
    2 = ALL PARAMETERS
    3 = PARAMETERS FROM EXISTING INDEXED DATA SET
ENTER OPTION: 1
```

[2] Respond to this prompt by entering the digit '1'. This response causes a one line prompt from the next prompt sequence to be displayed.

Note: Although the following prompts are displayed one line at a time when using the utility, all the prompts are listed here in logical groups for simplicity in describing the parameters.

```
[3] SECONDARY INDEX (Y/N)? : Y
[4] ENTER SECONDARY DATASET (NAME,VOLUME): CITY,EDX002
[5] SECONDARY KEY SIZE :4
[6] SECONDARY KEY POSITION :5

DATA SET SIZE IN EDX RECORDS: 10
INDEXED ACCESS METHOD RETURN CODE: -1
SYSTEM RETURN CODE: -1
CREATE/DEFINE FILE (Y/N)? : Y
NEW DATA SET IS ALLOCATED
DO YOU WANT IMMEDIATE WRITE-BACK? N
INVOKE LOAD(L), REORGANIZE(R) OR END(E) AFTER CURRENT FUNCTION? L
DEFINE IN PROGRESS
DATA SET SIZE IN EDX RECORDS: 10
INDEXED ACCESS METHOD RETURN CODE: -1
SYSTEM RETURN CODE: -1
PROCEED WITH LOAD/REORGANIZE (Y/N)? Y
[7] SECONDARY INDEX LOAD ACTIVE
ANSWER NULL FOR ALLOCATING DEFAULT WORK DATASET $SORTWRK
[8] SORT WORK DATASET (DSNAME,VOLUME):
SORT WORK DATASET REQUIRES 20 EDX RECORDS
5 RECORDS LOADED
SECONDARY LOAD SUCCESSFUL
```

- [3] Reply Y to this prompt because you are defining a secondary index.
- [4] Enter the data set name and volume where this index is being defined.
- [5] Specify the length of the secondary key within the primary index record for which this index is being defined.
- [6] Specify the starting position of the secondary within the primary index file record. The first byte of the record is number 1.
- [7] The secondary index load function is active.
- [8] At this point there are three possible responses:
 1. A null response, just pressing the enter key, will cause \$SORTWRK to be allocated on the IPL volume, if space is available. The size of the data set is calculated by the utility. If \$SORTWRK already exists, this indicates that another user is using the default work data set and you will be prompted again for a work data set name.

2. Entering a comma (,) followed immediately with a volume name, then pressing the enter key, causes the utility to try to allocate \$SORTWRK on the specified volume.
3. Entering a data set name and optionally a volume name (no volume name entered causes the IPL volume to be used) causes the utility to calculate the size of data set required and allocate it according to your response.

Notes:

1. If \$IAMUT1 allocates the data set for you, the data set will be automatically deleted at the end of the sort operation. However, if you provide either the name of an already existing data set (other than \$SORTWRK) or a data set name you want \$IAMUT1 to allocate, the data set will not be deleted at the end of the sort.
2. The sort work data set cannot always be calculated precisely because the size is dependent on several variables related to the input file. In most cases the calculated size will be adequate. However, if the size calculated is too small, the sort will end prematurely. If this happens you can preallocate a data set with a larger size than that calculated by \$IAMUT1 and execute the sort again.

OPTION 2

The \$IAMUT1 Indexed Access Method utility can be loaded with the Event Driven Executive operator command \$L \$IAMUT1.

When \$IAMUT1 is loaded the first prompt displayed is as follows:

```
[1] ENTER COMMAND (?): SE
```

[1] Entering SE causes the following option list prompt to be displayed.

```
SET FILE DEFINITION PARAMETERS
0 = EXIT
1 = SIGNIFICANT PARAMETERS
[2] 2 = ALL PARAMETERS
3 = PARAMETERS FROM EXISTING INDEXED DATA SET
ENTER OPTION: 2
```

[2] Respond to this prompt by entering the digit '2'. This response causes a one line prompt from the next prompt sequence to be displayed.

Note: Although the following prompts are displayed one line at a time when using the utility, the entire prompt list is shown for simplicity in presentation.

```
[3] SECONDARY INDEX (Y/N)? : Y
[4] ENTER SECONDARY DATASET (NAME,VOLUME): NAME,EDX002
[5] PARAMETER  DEFAULT NEW VALUE
    BASEREC      20:20
    BLKSIZE      256:
    KEYSIZE      4:
    KEYPOS       5:
    FREEREC      0:
    FREEBLK      0:
    RSVBLK       NULL:
    RSVIX        0:
    FPOOL        NULL:
    DELTHR       NULL:
    DYN          NULL:05
    TOTAL LOGICAL RECORDS/DATA BLOCK:      15
    FULL RECORDS/DATA BLOCK:               15
    INITIAL ALLOCATED DATA BLOCKS:        2
    INDEX ENTRY SIZE:                      12
    TOTAL ENTRIES/INDEX BLOCK:            20
    FREE ENTRIES/PIXB:                     0
    RESERVE ENTRIES/PIXB(BLOCKS):         0
    FULL ENTRIES/PIXB:                     20
    RESERVE ENTRIES/SIXB:                  0
    FULL ENTRIES/SIXB:                     20
    DELETE THRESHOLD ENTRIES               0
    FREE POOL SIZE IN BLOCKS                5
    # OF INDEX BLOCKS AT LEVEL 1:          1
    DATA SET SIZE IN EDX RECORDS:         10
    INDEXED ACCESS METHOD RETURN CODE:      -1
    SYSTEM RETURN CODE:                    -1
    CREATE/DEFINE FILE (Y/N)? : Y
    DATA SET EXISTS ALREADY AND IS LARGE ENOUGH
    DO YOU WISH TO USE IT (Y/N): Y
    DO YOU WANT IMMEDIATE WRITE-BACK? N
    INVOKE LOAD(L), REORGANIZE(R) OR END(E) AFTER CURRENT FUNCTION? L
    DEFINE IN PROGRESS
    DATA SET SIZE IN EDX RECORDS:         10
    INDEXED ACCESS METHOD RETURN CODE:      -1
    SYSTEM RETURN CODE:                    -1
    PROCEED WITH LOAD/REORGANIZE (Y/N)? Y
[6] SECONDARY INDEX LOAD ACTIVE
    ANSWER NULL FOR ALLOCATING DEFAULT WORK DATASET $SORTWRK
[7] SORT WORK DATASET (DSNAME,VOLUME):
    SORT WORK DATASET REQUIRES      20 EDX RECORDS
        5 RECORDS LOADED
    SECONDARY LOAD SUCCESSFUL
```

[3] Reply Y to this prompt because you are defining a secondary index.

[4] Enter the data set name and volume where this index is being defined.

[5] The following parameter list allows you to precisely define the secondary index structure.

[6] The secondary index load function is active.

[7] If the name of a data set and volume are entered, Sort/Merge will use it for the work data set. If a null response is made, the utility will calculate the size data set required and allocate it with the name \$SORTWRK on the IPL volume.

Notes:

1. For a more complete description of the responses available and the possible conditions that could exist, see step 8 description under "Option 1."
2. The following messages are from the IBM Sort/Merge Program Product, program number 5719-SM2. The following message list is the result of the secondary load function calling and executing Sort/Merge. For a description of the Sort/Merge program and its messages refer to IBM Series/1 Event Driven Executive Sort/Merge Programmer's Guide, SL23-0016..

```
SORT099N ----+----1----+----2----+----3----+----4----+----5----+
SORT000* LO$SYSPRTR
SORT001P SORT/MERGE SPECIFICATION PHASE STARTED
SORT000* HSORTR 12A0DP
SORT000* DW $SORTWRKTVOL
SORT000* FNC0001 12
SORT000* FR
SORT075P SPECIFICATION PHASE ENDED
SORT076P INPUT PHASE STARTED
SORT082P INPUT PHASE ENDED          5          1          1
SORT085P FINAL MERGE PHASE STARTED          1          4          5
SORT086P FINAL MERGE PHASE ENDED
SORT088P RECORDS READ FROM INPUT DATA SET(S):          0
SORT089N RECORDS INSERTED BY INPUT EXIT ROUTINE:          5
SORT090N RECORDS DELETED BY INPUT EXIT ROUTINE:          0
SORT091N RECORDS INSERTED BY OUTPUT EXIT ROUTINE:          0
SORT092N RECORDS DELETED BY OUTPUT EXIT ROUTINE:          0
SORT093N RECORDS WRITTEN TO OUTPUT DATA SET:          0
SORT094N I/O ERRORS ACCEPTED:          0
SORT095N I/O ERRORS SKIPPED:          0
SORT149N RECORDS SORTED OR MERGED:          5
SORT097P NORMAL ENDING FOR SORT/MERGE PROCESSING
```

OPTION 3

The \$IAMUT1 Indexed Access Method utility can be loaded with the Event Driven Executive operator command \$L \$IAMUT1.

When \$IAMUT1 is loaded the first prompt displayed is as follows:

```
[1] ENTER COMMAND (?): SE
```

[1] Entering SE causes the following option list prompt to be displayed.

```
SET FILE DEFINITION PARAMETERS
0 = EXIT
1 = SIGNIFICANT PARAMETERS
2 = ALL PARAMETERS
[2] 3 = PARAMETERS FROM EXISTING INDEXED DATA SET
ENTER OPTION: 3
```

[2] Respond to this prompt by entering the digit '3'. This response causes a one line prompt from the next prompt sequence to be displayed.

Note: Although the following prompts are displayed one line at a time when using the utility, the entire prompt list is shown for simplicity in describing the parameters.

```
[3] SECONDARY INDEX (Y/N)? : Y
[4] ENTER SECONDARY DATASET (NAME,VOLUME): LEVEL,EDX002
[5] NAME OF EXISTING INDEXED DATA SET (NAME,VOLUME): CITY,EDX002
[6] NEW PARAMETERS EXACTLY SAME AS ORIGINAL PARAMETERS (Y/N)? Y
    DATA SET SIZE IN EDX RECORDS:          10
    INDEXED ACCESS METHOD RETURN CODE:      -1
    SYSTEM RETURN CODE:                    -1
    CREATE/DEFINE FILE (Y/N)? : Y
    NEW DATASET IS ALLOCATED
    DO YOU WANT IMMEDIATE WRITE-BACK? Y
    INVOKE LOAD(L), REORGANIZE(R) OR END(E) AFTER CURRENT FUNCTION? L
    DEFINE IN PROGRESS
    DATA SET SIZE IN EDX RECORDS:          10
    INDEXED ACCESS METHOD RETURN CODE:      -1
    SYSTEM RETURN CODE:                    -1
    PROCEED WITH LOAD/REORGANIZE (Y/N)? Y
    SECONDARY INDEX LOAD ACTIVE
    ANSWER NULL FOR ALLOCATING DEFAULT WORK DATASET $SORTWRK
[7] SORT WORK DATASET (DSNAME,VOLUME):
    SORT WORK DATASET REQUIRES          20 RECORDS
        5 RECORDS LOADED
    SECONDARY LOAD SUCCESSFUL
```

[3] Reply Y to this prompt because you are defining a secondary index.

[4] Enter the data set name and volume where this index is being defined.

[5] Enter the data set name and volume of the secondary index whose parameters are to be used for this index.

[6] If all parameters are to be the same as those initially set for the data set name entered in prompt **[3]**, reply Y. However, if you want the parameters adjusted, based on current file status, reply N.

[7] If the name of a data set and volume are entered, Sort/Merge will use it for the work data set. If a null response is made, the utility will calculate the size data set required and allocate it with the name \$SORTWRK on the IPL volume.

Note: For a more complete description of the responses available and the possible conditions that could exist, see step 8 description under "Option 1."

LOADING A SECONDARY FILE WITH AN APPLICATION PROGRAM

You have the option of allowing \$IAMUT1 to load your secondary file at the time it is created, as was demonstrated in "Option 2" on page IG-79, or you can load it with an application program. The sequence of operation for loading your secondary index with an application program is described here.

A secondary file has the following format:

Secondary Key	Sequence Number	Primary Key	Relative Block Number
SMITH	0001	12345AB	RBN

┌─── 4 bytes ──┐ ┌─── 4 bytes ──┐

In preparation for loading your secondary index, allocate the following sort data sets. The size of the input and output data sets are calculated using the lengths of keys from the primary index file record plus four bytes for the sequence number and four bytes for the RBN.


Secondary key length + primary key length + 8 bytes

- Sort input data set
- Sort output data set
- Sort work data set

If you have the Sort/Merge licensed program product, program number 5719-SM2, refer to the IBM Series/1 Event Driven Executive Sort/Merge Programmer's Guide, SL23-0016, for details of the sort work data set and sort specification data sets. Otherwise, use your own sort program.

Loading your secondary index requires the following sequence:

1. Open the primary index file in process mode
2. Retrieve a primary index file record with a GETSEQR request (retrieves the primary record plus the RBN)
 - a. When end-of-data condition is reached, go to step 6
3. Using values from the retrieved record, build a secondary record with the following format:



4. Move the newly built secondary record into the sort input data set.
5. Return to step 2 and repeat the sequence.
6. Sort the records in the sort input data set using the following sort specifications.
 - a. Sort the input records on position 1 through n-4 into ascending sequence (n = the length of the records as calculated previously for the sort data sets).
7. Open the secondary index, which is to be loaded, in load mode.

If your program is written in Event Driven Language (EDL), specify the independent option when you open the secondary index. If you are using a high level language, use \$IAMUT1 to turn on the independent indicator.
8. Read a record from the sorted output.
 - a. When end-of-file is reached, go to step 14.
9. Move a sequence number into the retrieved record's sequence number field (use X'0000' for the first record).
10. Increment the sequence number by a value of 1.
11. Use a PUT request to load the record into the secondary file.
12. Return to step 8 and repeat the sequence.
13. Issue a DISCONNECT to the primary index file and secondary index.
14. Using the \$IAMUT1 utility, turn off the invalid indicator for this secondary index entry in your directory. Also, turn off the independent indicator if you turned it on in step 7.



CHAPTER 6. PROCESSING THE INDEXED FILE

This chapter provides information for designing applications that use the Indexed Access Method. It contains information about:

- Connecting and disconnecting the indexed file
- Accessing the indexed file
- Maintaining the indexed file

"Chapter 7. Coding the Indexed Access Method Requests" contains a detailed description of the EDL coding syntax of each Indexed Access Method request. You may wish to refer to it while reading this chapter.

CONNECTING AND DISCONNECTING THE INDEXED FILE

An indexed file must be defined and formatted by using the \$IAMUT1 utility set parms (SE) and define (DF) commands before issuing a LOAD or PROCESS request to the file.

Prior to using an indexed file, you must issue either a LOAD or PROCESS request to connect it to your program. The file must be defined in your PROGRAM statement or by a DSCB statement. A CALL statement specifying either LOAD or PROCESS automatically opens the Indexed Access Method file. If you have an already open DSCB for the Indexed Access Method file you can pass it as a parameter, but that is not required.

CONNECTING

A LOAD or PROCESS request builds an indexed access control block (IACB) that is associated with an indexed file. The IACB connects a request to the file.

In load mode, data records are placed in the file sequentially (free records and blocks are skipped). When in process mode, data records are placed in the first appropriate slot in the file (free space is used) unless the record has a new high key. In the case of a new high key, the record is placed in the next available base record slot.

Only one LOAD request can be active for a given file. However, processing can take place concurrently with loading.

Multiple IACBs can be associated with the same file. Data integrity is maintained by a locking system that assigns file locks, record locks, or block locks to the requesting IACB. This prevents concurrent modification of index or data records, thereby avoiding the possibility of a double update situation.

Some applications will not need to wait for a lock to be released on a record, block, or buffer. In these situations you might want to use the conditional requests available for some Indexed Access Method functions. The conditional function requests allow control to be returned immediately to the requesting program for other processing, then return later to attempt to retrieve the record which was locked. The conditional requests are described in "Chapter 7. Coding the Indexed Access Method Requests."

An IACB can hold only one lock at a time; if your application requires concurrent execution of functions that obtain locks (direct update or sequential update - see "Accessing the Indexed File" on page IG-89 for a description of these functions), you must issue multiple PROCESS requests to build multiple IACBs.

DISCONNECTING

A DISCONN request disconnects an IACB from the file, releases the storage for that IACB, releases locked blocks or records being held by that IACB, and writes out to disk any blocks that are being held in the buffer. The DISCONN request can be issued at any time during loading or processing.

There is no automatic DISCONN on task termination. Failure to disconnect your indexed files prior to task termination may prevent resources that were allocated to your task from being allocated to other tasks and updated records from being written to your file.

Using Secondary Keys

To access a file by a secondary key, you issue either a LOAD or PROCESS request, specifying the file name of the **secondary index** and specifying **secondary keys** when referencing data records. The Indexed Access Method determines the relationships among the files by using the directory and automatically opens the primary file. All subsequent operations done under this LOAD or PROCESS access the file using the secondary index. You must open a file by the primary name to access it by the primary keys.

Direct retrievals use the secondary index, and sequential retrievals return records in order by secondary key. Records within a group which have the same secondary key are returned in the order which the records were written into the file. Each application must determine whether the correct record has been retrieved when duplicate keys are possible; the Indexed Access Method provides no facility for that determination.

When records are updated, inserted, or deleted, in primary index files, some or all secondary indexes can be updated automatically according to the options you selected in the directory entries. These options are: auto-update and independent processing.

If the auto-update indicator is on in the directory entry for a secondary index and you open the associated primary file to insert, delete, or update records, the associated secondary index will be updated automatically. There is no consideration for whether the independent indicator

is on for the secondary. However, if the invalid indicator for the secondary entry is on, the secondary index is not updated.

The independent indicator is used when a secondary index is opened in load mode to add new entries to the file.

Note: When records are accessed by their secondary key, you must ascertain through your application program that you have retrieved the correct record because of the possibility of duplicate keys.

ACCESSING THE INDEXED FILE

Initiate general purpose access to an indexed file with a PROCESS request. After the PROCESS request has been issued, any of the following functions can be requested:

- Direct reading - Retrieving a single record independently of any previous request.
- Direct updating - Retrieving a single record for update; complete the update by either replacing, deleting, or releasing the record.
- Sequential reading - Retrieving the next logical record relative to the previous sequential request.

The first sequential request can access the first record in the file or any other record in the file by key (except COBOL applications).

- Sequential updating - Retrieving the next logical record for update; complete the update by either replacing, deleting, or releasing the record.
- Inserting - Placing a single record, in its logical key sequence, into the indexed file.
- Deleting - Removing a single record from the indexed file.
- Extracting - Extracting data that describes the file.

Note that the update functions require more than one request.

When a function is complete, another function may be requested, except that a sequential processing function may be followed only by another sequential function. You can terminate sequential processing at any time by issuing a DISCONN or ENDSEQ request. An end-of-data condition also terminates sequential processing.

DIRECT READING

Use the GET request to read a record using direct access. The key parameter is required and must be the address of a field of full key length regardless of the key length specification.

The record retrieved is the first record in the file that satisfies the search argument defined by the key and key relation (krel) parameters. The key field in your program is updated to reflect the key contained in the record that satisfied the search.

If the key length is specified as less than the full key length, only part of the key field is used for comparison when searching the file. For example, the keys in a file are AAA, AAB, ABA, and ABB, the key field contains AB0, and key relation is EQ. If key length is zero, the search argument defaults to the full key AB0 and a record-not-found code is returned. If the key length specification is 2 and the search argument is AB, the third record is returned. If the key length specification is 1 and the search argument is A, the first record is returned.

DIRECT UPDATING

To update a record using direct access:

1. Retrieve the record with a GET request, specifying the key and key relation (krel) parameters.
2. Complete the update by doing one of the following:
 - If you want to change the record, modify the record in your buffer (do not change the key field of the record). Issue a PUTUP request to return the updated record to the file.
 - If you do not want to change the record, issue a RELEASE request.
 - If you want to delete the record, issue a PUTDE request.

The key parameter must be specified as the address of a field of full key length. The primary key cannot be modified during the update; a secondary key can.

The only valid requests, other than DISCONN and EXTRACT, that can follow GET for direct update are PUTUP, PUTDE, and RELEASE.

During the update, the subject record is locked (made unavailable) to any other request until the update is complete. Even if no action is taken after the GET request is issued, the RELEASE request is required to release the lock on the record. You may wish to use the conditional option on your requests to avoid unnecessary wait for locks. For conditional request coding see "Chapter 7. Coding the Indexed Access Method Requests" on page IG-97. For details on long lock time or dead lock condition, see "Deadlocks and the Long-Lock-Time Condition" on page IG-218.

SEQUENTIAL READING

Use the GETSEQ request for sequential access to records. After a sequential processing request has been initiated, only sequential functions can

be requested until an end-of-data condition occurs or an ENDSEQ request is issued. Processing is terminated when a DISCONN request is issued or an error or warning is returned.

Figure 11 on page IG-92 summarizes the protocol for sequential processing.

Note: You can sequentially process a file more than once.

To begin sequential access with the first record in a file, set the key address to zero. To start with any other record, specify a search argument by specifying the key and key relation (krel) parameters.

If you specify a search argument, the key field is modified to reflect the key of the first record found.

After the first retrieval, a GETSEQ retrieves the next sequential record regardless of any key or key relation specification. Therefore, you can use the same GETSEQ statement to retrieve all records. A search argument on succeeding retrievals is ignored and the key field is not modified.

When using secondary keys, you access the duplicate keys with a sequential get request. For example:

```
GETSEQ SMITH
```

Issuing the same request repeatedly will return all of the secondary keys whose value is SMITH. You must check to determine when the key changes, or when you have retrieved the particular record you want within that sequence of keys.

Specify ENDSEQ to stop reading before the end of data is reached. Reading ends automatically at the end of data. The end-of-data condition occurs when an attempt is made to retrieve a record after the last record in the file.

If you specify the end-of-data exit (EODEXIT) parameter on the PROCESS request, control is transferred to the address specified by the EODEXIT parameter when the end-of-data condition occurs.

During sequential reading, the block that contains the record is locked, making all records in the block unavailable to other requesters until the last record of the block is processed or sequential processing is ended. For details on long lock time or dead lock condition, see "Deadlocks and the Long-Lock-Time Condition" on page IG-218.

SEQUENTIAL UPDATING

To update a record using sequential access:

1. Retrieve the record with a GETSEQ request for update, specifying the key and one of the update key relation (krel) parameters. The key is used only on the first retrieval. Do not specify a key if processing is to begin with the first record in the file.

2. Complete the update by doing one of the following:

- If you want to change the record, modify the record in your buffer (do not attempt to change the primary key field of the record). Issue a PUTUP request to return the updated record to the file.
- If you do not want to change the record, issue a RELEASE request.
- If you want to delete the record, issue a PUTDE request.

During sequential updating, the block that contains the record is locked, making all records in the block unavailable to other requesters until the last record of the block is processed or sequential processing is ended.

Terminate processing with an ENDSEQ request or a DISCONN request either before or after completing the update. Processing is also terminated on an end-of-data condition.

REQUEST/CONDITION	CAN BE FOLLOWED BY:
GETSEQ	DISCONN END-OF-DATA CONDITION ENDSEQ PUTUP PUTDE RELEASE
End-of-data condition or ENDSEQ request	DISCONN GET PUT DELETE
PUTUP	DISCONN ENDSEQ GETSEQ
PUTDE	DISCONN ENDSEQ GETSEQ
RELEASE	DISCONN ENDSEQ GETSEQ

Figure 11. Protocol for Sequential Updating

INSERTING RECORDS

To insert a new record in a file, issue a PUT request after the file has been connected with a PROCESS. The Indexed Access Method uses the primary key of the record to insert the record into the file.

The primary key of the inserted record must be different from any key in the file; otherwise, a duplicate key error occurs. The key can be higher than any key in the file.

If you are not loading base records, and want to insert records into the file in random order, the following should be satisfied:

- for files defined by option 1, ensure that random (R) was specified
- for files defined by option 2, ensure that sufficient free pool space was specified

DELETING RECORDS

Use DELETE to delete a record from the file. Specify the full key of the record. If no record exists with the specified key, a warning return code is returned.

EXTRACTING INDEXED FILE INFORMATION

The EXTRACT request provides information about a file from the file control block (FCB) or FCB Extension. It can also return data paging statistics to the calling program with an option to reset the counters. Data paging is described under "Data Paging" on page IG-206.

The FCB includes information such as key length, key displacement, block size, record size, and other data regarding the file structure. The FCB Extension contains the \$IAMUT1 utility SE command parameters that were used to define the file.

The EXTRACT request copies the file control block or the FCB Extension to an area that you provide. The file must have been connected by a LOAD or PROCESS request.

The contents of the FCB block and the FCB Extension are described by FCBEQU, a unit of copy code that is supplied with the Indexed Access Method. Use COPY FCBEQU to include these equates in an EDL program.

An EXTRACT issued for a secondary file returns the primary FCB with the secondary key size and position of the secondary key. If you want the actual FCB of the secondary file, you must open the secondary file independent and then the secondary index FCB will be returned for the EXTRACT request. The FCB extension returned is always the secondary FCB extension.

MAINTAINING THE INDEXED FILE

This section describes how to maintain Indexed Access Method files. The following topics are discussed:

- File backup and recovery
- File recovery without backup
- Reorganizing the file
- Dumping the file
- Deleting the file
- Verifying an indexed file

FILE BACKUP AND RECOVERY

To protect against the destruction of data, copy the indexed file (or the volume in which the file exists) at regular intervals using the \$COPY utility. (See the Operator's Reference for instructions on using the Event Driven Executive utilities.) To obtain a sequential dump of an indexed file, use the \$IAMUT1 utility UN command. During the interval between making copies, you should keep a journal file of all transactions made against the indexed file.

The journal file can be a consecutive file containing records that describe the type of transaction and the pertinent data. A damaged indexed file can be recovered by updating the backup copy from the journal file.

For example, suppose an indexed file named REPORT is lost because of physical damage to the disk. The condition that caused the error has been repaired and the file must be recovered. Delete REPORT, copy the backup version of REPORT to the desired volume, and process the journal file to recreate the file.

If a data-set-shut-down condition exists, cancel \$IAM and reload it. Then issue a PROCESS to the REPORT file and, using the journal file, reprocess the transactions that occurred after the backup copy was made. Refer to "The Data-Set-Shut-Down Condition" on page IG-215 for more information on the data-set-shut-down condition.

Backing Up A Secondary Index

A Secondary index can be backed up the same as primary index files. However, if your primary file is backed up you can rebuild your secondary from the backup copy of the primary indexed file.

Duplicate secondary keys are maintained in the order they are inserted by a secondary key sequence number. This sequence number is incremented with

each new insert. When a secondary index is reloaded the secondary key sequence numbers are reassigned. Therefore, the history of which records were written to the file first is lost.

Note: If your application is dependent on the secondary key sequence number history, you would not want to rebuild your secondary index because the sequence numbers are reset.

RECOVERY WITHOUT BACKUP

If you do not use the backup procedures as described previously under "File Backup and Recovery" on page IG-94, and you encounter a problem with your file, you still may be able to recreate your file. However, the status of requests that were in process at the time of the problem is uncertain.

To recreate your file, follow the steps in "Reorganizing an Indexed File" to reorganize your file. After recreating the file, verify the status of the requests that were in process when the problem occurred.

REORGANIZING AN INDEXED FILE

An indexed file must be reorganized when a record cannot be inserted because of lack of space. This condition does not necessarily mean that there is no more space in the file; it means that there is no space in the area where the record would have been placed. Therefore, you may be able to reorganize without increasing the size of the file. Perform the following steps to reorganize a file:

1. Ensure that all outstanding requests against the file have been completed; issue a DISCONN for every current IACB.
2. Use the set parms (SE) or define (DF) commands of the \$IAMUT1 utility to define a new indexed file. Estimate the number of base records and the amount and mix of free space in order to minimize the need for future reorganizations. See "Chapter 3. Defining Primary Index Files" for guidelines for making these estimates.

You can use Option 3 of the SE command to define the new file like the original indexed file.

3. Use the reorganize command (RO) of the \$IAMUT1 utility to load the new indexed file from the indexed file to be reorganized.

Alternatively, you can use the unload command (UN) of the \$IAMUT1 utility to transfer the data from an indexed file to a sequential file, then use the load command (LO) to load it back into the indexed file.

4. Use the \$DISKUT1 utility to delete the old file and rename the new file.

REORGANIZING A SECONDARY INDEX: Reorganizing a secondary index does not reset the secondary key sequence numbers during the reorganization. The records are placed in another Indexed Access Method file without any modification within the individual records. The secondary key sequence numbers will be reset however, when the index is loaded.

DUMPING AN INDEXED FILE

To produce a hexadecimal dump of an indexed file, use the DP command of the \$DISKUT2 utility. The dump includes control information, index blocks, and data blocks. For information on the \$DISKUT2 utility, refer to the Operator's Reference. The Indexed Access Method Version 2 Internal Design contains information on dumping an indexed file.

DELETING AN INDEXED FILE

Delete an indexed file the same way you delete any other file. From a terminal, use the DE command of the \$DISKUT1 utility; from a program, use the \$DISKUT3 data management utility. (Refer to the Operator's Reference for a description of \$DISKUT1, and to the System Guide for a description of \$DISKUT3).

VERIFYING AN INDEXED FILE

\$VERIFY helps you check the validity of an indexed file and prints control block and free space information about the file on \$SYSPRTR.

With \$VERIFY you can:

- Verify that all pointers in an indexed file are valid and that the records are in ascending sequence by key.
- Print a formatted File Control Block (FCB) listing, including the FCB Extension block. The FCB Extension block contains the original file definition parameters.
- Print a report showing the distribution of free space in your file.
- Verify secondary files against primary files.

For details on using \$VERIFY, see "Chapter 9. The \$VERIFY Utility" on page IG-189.

CHAPTER 7. CODING THE INDEXED ACCESS METHOD REQUESTS

This chapter describes the syntax used to code Event Driven Language requests for the Indexed Access Method.

The information in this chapter is intended for use as a reference when coding EDL application programs that use the Indexed Access Method. For information on coding Indexed Access Method applications in other languages, refer to the appropriate language manual.

Included for each request is a description of the purpose of the request, the detailed coding syntax, a description of each parameter, and all of the return codes associated with using these requests.

At the end of this chapter is a summary of the syntax of the EDL CALL instructions used to invoke the functions provided by the Indexed Access Method.

For a complete example of using the Indexed Access Method requests, refer to "Appendix C. Coding Examples" on page IG-233.

REQUEST FUNCTIONS OVERVIEW

This section provides an overview of the Indexed Access Method requests and how to code them. The Indexed Access Method callable requests are:

<u>Request</u>	<u>Description</u>
DELETE	Deletes a single record, identified by its key, from the file. Use DELETE to delete a record; the record cannot have been retrieved for update.
DISCONN	Disconnects an IACB from an indexed file, thereby releasing any locks held by that IACB; writes out all buffers associated with the file; and releases the storage used by the IACB.
ENDSEQ	Terminates sequential processing.
EXTRACT	Provides information about the file from the File Control Block, File Control Block Extension, and data paging statistics.
GET	Directly retrieves a single record from the file. If you specify the update mode, the record is locked (made unavailable to other requests) and held for possible modification or deletion. Use GET to retrieve a single record from the file.
GETSEQ	Sequentially retrieves a single record from the file. If you specify update mode, the block containing the record is locked

(made unavailable to other requests) and held for possible modification or deletion. Use GETSEQ when you are performing sequential operations.

LOAD Builds an Indexed Access Control Block (IACB) and connects it to an indexed file. You can then use the IACB to issue LOAD requests to that file to load records.

PROCESS Builds an Indexed Access Control Block (IACB) and connects it to an indexed file. You can then use the IACB to issue requests to that file to read, update, insert, and delete records. A program can issue multiple PROCESS functions to obtain more than one IACB for the same file, enabling the file to be accessed by several requests concurrently within the same program.

PUT Loads or inserts a new record depending on whether the file was opened with the LOAD or PROCESS request. Use PUT when you are adding records to a file.

PUTDE Deletes a record that is being held for update. Use PUTDE to delete a record that has been retrieved in update mode.

PUTUP Replaces a record that is being held for update. Use PUTUP to modify a record.

RELEASE Releases a record that is being held for update. Use RELEASE when a record that was retrieved for update is not changed.

CODING INDEXED ACCESS METHOD REQUESTS

All Indexed Access Method services are requested by using the CALL instruction. Parameters on the CALL instructions can have the following forms:

NAME: passes the value of the variable with the label 'NAME'

(NAME): passes the address of the variable 'NAME' or the value of a symbol defined using an EQU statement

For additional information, refer to the description of the CALL instruction in the Language Reference.

General Statement Format

The general form of all Indexed Access Method calls is as follows:

```
CALL IAM,(func),iacb,(parm3),(parm4),(parm5)
```

The request type is determined by the operand 'func'. In addition to the function request, you will notice that some functions allow a suffix of C, R, or CR. The C means perform the function requested conditionally. The condition is that the function is to be executed only if the record, the block containing the record, or the buffer containing the record is not locked. If any of those three items are locked for the record being requested, control is to be returned to the requesting program immediately. A return code is set to indicate that a lock was encountered. A conditional request can still wait on a resource if it is during the process of updating an index for a delete or insert.

The appended character, R, means return the record and the relative block number (RBN) of the record being requested. Again this can be a conditional request by preceding the letter R with the letter C. The combination CR, indicates that the record and RBN is to be returned conditionally; return the record and RBN only if the record, block, or buffer is not locked by another request.

If the RBN is requested and the record, block, and buffer are free, the RBN is returned as a 4-byte value. The 4-byte RBN value is returned at the end of the retrieved record. Therefore, when using the suffix R, ensure that your buffer is large enough to accommodate the record length, plus the 4-byte RBN value.

The RBN can be used if you are building or maintaining your own secondary index. However, because records in an indexed file are subject to being moved to different locations (RBNs) due to insert and delete activity, the RBN is not guaranteed to remain accurate if insert and delete activity to the primary index file occur.

The option of C, R, or CR is indicated in the boxed instructions with a vertical bar (|). The presence of this bar indicates that a choice must be made. Only one of the requests can be used in any one statement. For example, PUT|PUTC, you must choose one or the other when coding the request.

Depending on the type of function the remaining parameters may or may not be required. The symbols used for func and parm5 are provided by EQU statements in the IAMEQU copy code module and are coded as shown in the syntax descriptions. These symbols are treated as addresses; therefore the MOVEA instruction should be used if it is necessary to move them into a parameter list. Since these symbols are equated to constants, they may also be manipulated using other instructions by prefixing them with a plus (+) sign. Use the COPY statement to include IAMEQU in your program.

Note: You can not use the software registers (#1 and #2) on Indexed Access Method calls.

Using Program Variables

If you use variables for parameters parm3, parm4, and parm5 (that is, you code them without parentheses or a plus sign), they are set to zero by the Indexed Access Method before returning. Those parameters must be reinitialized before executing the CALL instruction again.

Link-edit Considerations

Programs which call the Indexed Access Method must be processed by \$EDX-LINK to include the subroutine module IAM. IAMEQU has an EXTRN statement for IAM. Refer to the chapter on "\$EDXLINK" in the Program Preparation Guide for information on how to perform the link-edit process.

Return Codes

All Indexed Access Method requests pass a return code reflecting a condition that prevailed when the request completed. This code is passed in the task code word (referred to by task name) of the TCB associated with the requesting task. These return codes fall into three categories:

-1	= Successful completion
Positive	= Error
Negative	= Warning (other than -1)

Note: Return codes 1, 7, 8, and 22 are positive value return codes but they do not cause the error exit routine to be entered, even when ERREXIT is coded. Also the negative (warning) return codes do not cause error exits. For details on coding ERREXIT, see "LOAD - Open

File for Record Loading" on page IG-118, or "PROCESS - Open File" on page IG-122.

The return codes associated with each request are included with the description of the request.

The Indexed Access Method also has the capability of logging errors in the system error log. Automatic updates for secondary indexes could encounter several errors within one request. These errors will be logged in the system error log if \$LOG is active. This may provide additional information when analyzing errors. For a description of \$LOG, refer to the System Guide.

CALL FUNCTION DESCRIPTIONS

The Indexed Access Method CALL functions are described on the following pages and are arranged in alphabetic order.

DELETE - DELETE RECORD

The DELETE request deletes a specific record from the file. The record to be deleted is identified by its key. The deletion makes space available for a future insert. The file must be opened in the PROCESS mode.

The DELETEDC request deletes a specific record from the file only if the record, block, or buffer is not locked.

Syntax:

label	CALL	IAM, (DELETE DELETEDC), iacb, (key)
Required:	all	
Defaults:	none	

Operands Description

iacb The label of a word containing the IACB address returned by PROCESS.

(key) The label of your key area containing the full key identifying the record to be deleted.

DELETE Return Codes

Code	Condition
-1	Successful
-85	Record not found
-90	Request cancelled because the request was conditional and a wait on a lock or buffer would be required
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
14	Invalid index block found - during processing an incorrect index block type was found, recreate the file
22	Address supplied by your program is not a valid IACB
76	DSOPEN error occurred - The system error field in the OPEN table contains the DSOPEN errors: 21 - DSNAME,VOLUME not found 22 - VOLSERV error 23 - I/O error
80	Write error - FCB. See system return code
100	Read error - check system return code
101	Write error - check system return code
230	Directory read error form \$IAMDIR - check system return code
242	Secondary index is out of sync with primary file. Must rebuild file to get back in sync
244	Error in opening auto-update on secondary modification request
245	Auto update PUTDE to a secondary failed, Auto-update processing continues
247	During auto-update processing a GETSEQ to a secondary failed, auto-update processing continues

DELETE Example

The following example deletes the record whose key is 'KEY0001' from the file. The file is identified by the field named 'FILE1'.

```
CALL  IAM,(DELETE),FILE1,(KEY)
.
.
.
FILE1 DATA F'0'      IACB ADDRESS FROM PROCESS
KEY   TEXT  'KEY0001',LENGTH=7
```

DISCONN - CLOSE FILE

The DISCONN request disconnects an IACB from an indexed file and releases the storage used for the IACB. It releases any locks held by that IACB and writes out any modified block from the file that are being held in the system buffer. Other users connected to this file are not affected.

Syntax:

```
label      CALL      IAM,(DISCONN),iacb
```

```
Required:  all  
Defaults:  none
```

Operands Description

iacb The label of a word containing the IACB address returned by PROCESS or LOAD.

DISCONN Return Codes

Code	Condition
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
12	Data set shut down due to error; see version 2 guide, error recovery
13	Module not included in load module
22	Address supplied by your program is not a valid IACB
100	Read error - check system return code
101	Write error - check system return code
110	Write error, file closed

DISCONN Example

The following example closes the file identified by the field named 'FILE1'.

```
CALL IAM,(DISCONN),FILE1
.
.
FILE1 DATA F'0' IACB ADDRESS FROM PROCESS
```

ENDSEQ - END SEQUENTIAL PROCESSING

The ENDSEQ request ends sequential processing, during which a block is locked and fixed in the system buffer. Sequential processing is normally terminated by an end-of-data condition. The ENDSEQ request is useful for freeing the locked block when the sequence need not be completed. ENDSEQ is valid only during sequential processing.

Note: After sequential processing has been terminated, it can be restarted again anywhere in the file.

Syntax:

```
label      CALL      IAM,(ENDSEQ),iacb
```

Required: all

Defaults: none

Operands Description

iacb The label of a word containing the IACB address returned by PROCESS.

ENDSEQ Return Codes

Code	Condition
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
22	Address supplied by your program is not a valid IACB

ENDSEQ Example

The following example ends sequential processing for the file identified by the field named 'FILE1'.

```
CALL  IAM,(ENDSEQ),FILE1
      .
      .
      .
FILE1 DATA F'0'    IACB ADDRESS FROM PROCESS
```


EXTRACT - GET FILE INFORMATION

The EXTRACT function returns information to the calling program. On a specific call, it performs one of the following:

- Returns information from a File Control Block (FCB). The FCB contains such things as the block size, key length, and data set and volume names of the indexed file. The FCBEQU copy code module contains a set of equates to map the File Control Block.

An EXTRACT request issued for a secondary file retruns the primary FCB with the secondary key size and key position for the secondary index. If you want the FCB of the secondary file, you must open the secondary index with the independent option then the secondary index FCB will be returned. The FCB extension returned is always the FCB extension for the secondary index.

- Returns information from a File Control Block Extension. The FCB Extension contains the parameters used to define the file. The FCBEQU copy code module contains a set of equates to map the FCB Extension.
- Returns data paging statistics. These can be used to calculate page "hit" ratios.
- Returns data paging statistics, then resets them to begin accumulating new statistics.

Syntax

```
label CALL IAM,(EXTRACT),iacb,(buff),(size),(type)
```

```
Required: iacb (only if type is FCBNRM or FCBEXT)  
buff
```

```
Defaults: size = Full FCB  
type = FCBNRM
```

Operands Description

iacb The label of a word containing the IACB address returned by process or load. Required only if type=FCBNRM or FCBEXT; otherwise ignored.

(buff) The label of the user area into which the data is returned.

If type=FCBNRM or FCBEXT, the File Control Block is returned in this area. The area must be large enough to contain the requested portion of the FCB. Use the COPY statement to include FCBEQU in your program so that the FCB and FCB Extension fields can be referenced by symbolic names.

If type=PAGST or PAGSTR, the paging statistics are returned in this area. In this case, the size parameter is ignored, and this area must be 16 bytes in length to accommodate the statistics. The paging statistics are returned in four double-word fields:

1. Write Miss Count
2. Write Hit Count
3. Read Miss Count
4. Read Hit Count

(size) Used only if type=FCBNRM or FCBEEXT; otherwise ignored. The number of bytes of the FCB or FCB Extension to be copied. The size of the FCB is the value of the symbol FCBSIZE in the FCBEQU equate table. The size of the FCB Extension is the value of the symbol FCBXSIZ in the FCBEQU equate table. Either of these symbols can be coded as the size parameter.

(type) Type of data to be returned. The following are defined:

FCBNRM Extract the FCB.

FCBEEXT Extract the FCB Extension.

PAGST Returns data paging statistics to the buffer. It always returns 16 bytes.

PAGSTR Same as PAGST, except the data paging statistics are reset to zero after being copied to the buffer. This allows a new set of statistics to be accumulated.

EXTRACT Return Codes

Code	Condition
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
22	Address supplied by your program is not a valid IACB
100	Read error - check system return code
120	Invalid extract type
122	File does not contain FCB extension
123	Cannot extract paging statistics. Data paging not active

EXTRACT Examples

The following example retrieves the current paging statistics and places them in the four double words provided.

```
CALL    IAM,(EXTRACT),0,(WRMIS),0,(PAGST)
      .
      .
      .
WRMIS   DATA  D'0'      WRITE MISS COUNT
WRHIT   DATA  D'0'      WRITE HIT COUNT
RDMIS   DATA  D'0'      READ MISS COUNT
RDHIT   DATA  D'0'      READ HIT COUNT
```

The following example gets the attributes of the file identified by the field named FILE1 from the FCB and places them in an area called WORK.

```
CALL    IAM,(EXTRACT),FILE1,(WORK),(FCBSIZE)
      .
      .
      .
FILE    DATA  D'0'      IACB ADDRESS FROM PROCESS
WORK    DATA  256F'0'   FCB COPY AREA
        COPY   FCBEQU    FCB EQUATES
```

GET - GET RECORD

The GET request retrieves a single record from the indexed file and places the record in a user area. The file must have been opened using the PROCESS request before the GET request is issued.

The requested record is located by key. The search may be modified by a key relation (krel) or a key length (klen). The first record in the file that satisfies the key condition is the one that is retrieved.

Retrieve for update can be specified if the requested record is intended for possible modification or deletion. The record is locked and remains unavailable to any other requests until the update is completed by a PUTUP, PUTDE or by a RELEASE. The record is also released if an error occurs or processing is ended with a DISCONN.

During an update, you must not change the primary key field in the record or the field addressed by the key parameter. The Indexed Access Method checks for and prohibits primary key modification.

The GETC request retrieves a single record from the indexed file and places the record in a user area only if the record, block, or buffer is not locked.

The GETR request retrieves the RBN of a specified record from the indexed file and places the record and RBN in a user area.

The GETCR request retrieves the RBN of a specified record from the indexed file and places the record and RBN in a user area only if the record, block, or buffer is not locked.

Syntax:

```
label      CALL  IAM,(GET|GETC|GETR|GETCR),iacb,(buff),(key),  
            (mode/krel)
```

Required: iacb,buff,key

Defaults: mode/krel=EQ

Operands

Description

iacb The label of a word containing the IACB address returned by PROCESS.

(buff) The label of the user area into which the requested record is placed. When the RBN is requested, the RBN is returned at the end of the record. The user buffer must be four bytes longer than the record length to accommodate the RBN.

(key)

The label of your key area containing the key identifying the record to be retrieved and preceded by the lengths of the key and area. This area has the standard TEXT format and may be declared using the TEXT statement. If you do not use the TEXT statement for this field, you must code it in the same format as the TEXT statement generates.

The TEXT statement format is as follows:

Offset	Field
key - 2	LENGTH (1 byte)
key - 1	KLEN (1 byte)
key	Key area ("LENGTH" bytes)

length The length of the key area. It must be equal to or greater than the full key length for the file in use.

klen The actual length of the key in the key area to be used as the search argument for the operation. It must be less than or equal to the full length of the keys in the file in use. If klen is 0, the full key length is assumed.

A generic key search is performed when klen is less than the full key size. The first n bytes (as specified by klen) of the key area are matched against the first n bytes of the keys in the file. The first matching key determines the record to be accessed. The full key of the record is returned in the key area.

key area The area containing the key to be used as a search argument. If you are using a generic key, after a successful GET request this area contains the full key of the record accessed.

(mode/krel) Retrieval type and key relational operator to be used. The following are defined:

EQ Retrieve only key equal

GT Retrieve only key greater than

GE Retrieve only key greater than or equal

UPEQ Retrieve for update key equal

UPGT Retrieve for update key greater than

UPGE Retrieve for update key greater than or equal

GET Return Codes

Code	Condition
-90	Request cancelled because the request was conditional and a wait on a lock or buffer would be required
-80	End of data
-58	Record not found
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
22	Address supplied by your program is not a valid IACB
100	Read error - check system return code
101	Write error - check system return code
200	Error occurred while accessing the primary file
242	Secondary index is out of sync with primary file.
247	During auto-update processing a GETSEQ to a secondary failed, auto-update processing continues.
248	I/O error on primary file during a secondary request.
249	GET UPDATE error occurred trying to update a bad RBN.

GET Example

The following example gets a record whose key is 'JONES'. The file records are 80 bytes in length and the key length is 20 bytes. The record is returned in the area named RECORD, and because this is a GETR request, the RBN is also returned in the area named RBN, which must follow immediately after the record area.

	CALL	IAM,(GETR),FILE1,(RECORD),(KEY)	
	.		
	.		
FILE1	DATA	F'0'	IACB ADDRESS FROM PROCESS
KEY	TEXT	'JONES',LENGTH=20	RECORD KEY
RECORD	DATA	128F'0'	RECORD AREA
RBN	DATA	D'0'	RBN

GETSEQ - GET RECORD (SEQUENTIAL MODE)

The GETSEQ request retrieves a single record from the indexed file and places the record in a user area (buff). The file must be opened in the PROCESS mode.

The first GETSEQ of a sequence is performed like a GET; the first record in the file that satisfies the key condition is the one that is retrieved. If key is zero, the first record in the file is retrieved. Subsequent requests in the sequence locate the next sequential record in the file and the key parameter is ignored if specified. The sequence is terminated by an end-of-data condition, by an ENDSEQ, by a DISCONN, or by an error. During the sequence, direct-access requests are invalid.

Retrieval for update can be specified if the requested record is intended for possible modification or deletion. If update is used the record is locked and remains unavailable to any other requests until the update is completed by a PUTUP, PUTDE or RELEASE. The record is also released by ending the sequence with an ENDSEQ or by ending processing with a DISCONN or by an error.

During an update, the user must not change the primary key field in the record or the field addressed by the primary key parameter. The Indexed Access Method checks for and prohibits key modification.

The GETSEQC request retrieves a single record from the indexed file and places the record in a user area only if the record, block, or buffer is not locked. The file must be opened in the PROCESS mode.

The GETSEQCR request retrieves the RBN of the specified record from the indexed file and places the record in a user area only if the record, block, or buffer is not locked. The file must be opened in the PROCESS mode.

Syntax:

label	CALL	IAM,(GETSEQ GETSEQC GETSEQR GETSEQCR),iacb, (buff),(key),(mode/krel)
-------	------	---

Required: iacb,buff,key

Defaults: mode/krel=EQ

Operand

Description

iacb The label of a word containing the IACB address returned by PROCESS.

(buff) The label of the user area into which the requested record is placed. When the RBN is requested, the RBN is returned at the

end of the record. The user buffer must be four bytes longer than the record length to accommodate the RBN.

(key)

The label of the user key area containing the key identifying the record to be retrieved and preceded by the lengths of the key and area. If the first record of the file is to be retrieved, this field as specified should be 0. The key field, if specified, has the standard TEXT format and may be declared using the TEXT statement. If you do not use the TEXT statement for this field, you must code it in the same format as the TEXT statement generates.

The TEXT statement format is as follows:

Offset	Field
key - 2	LENGTH (1 byte)
key - 1	KLEN (1 byte)
key	Key area ("LENGTH" bytes)

length The length of the key area. It must be equal to or greater than the full key length for the file in use.

klen The actual length of the key in the key area to be used as the search argument for the operation. It must be less than or equal to the full length of the keys in the file in use. If klen is 0, the full key length is assumed.

A generic key search is performed when klen is less than the full key size. The first n bytes (as specified by klen) of the key area are matched against the first n bytes of the keys in the file. The first matching key determines the record to be accessed. The full key of the record is returned in the key area.

key area The area containing the key to be used as a search argument. If you are using a generic key, after the first successful GETSEQ request this area contains the full key of the record accessed.

(mode/krel) Retrieval type and key relational operator to be used. The following are defined:

EQ Retrieve only key equal

GT Retrieve only key greater than

GE Retrieve only key greater than or equal

- UPEQ Retrieve for update key equal
- UPGT Retrieve for update key greater than
- UPGE Retrieve for update key greater than or equal

After the first GETSEQ of a sequence only the retrieval type is meaningful. The keys are not checked for equal or greater than relationship.

GETSEQ Return Codes

Code	Condition
-90	Request cancelled because the request was conditional and a wait on a lock or buffer would be required
-80	End of data
-58	Record not found
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
22	Address supplied by your program is not a valid IACB
100	Read error - check system return code
101	Write error - check system return code
200	Error occurred while accessing the primary file
242	Secondary index is out of sync with primary file.
248	I/O error on primary file during a secondary request.
249	GET UPDATE error occurred trying to update a bad RBN.

GETSEQ Example

The following example gets the record whose key is 'KEY0001' and places it in an area called 'BUFFER'. The file is identified by the field named 'FILE1'. Subsequent GETSEQ requests result in the next sequential record being returned.

```
CALL    IAM,(GETSEQ),FILE1,(BUFFER),(KEY)
      .
      .
      .
FILE1   DATA  F'0'           IACB ADDRESS FROM PROCESS
BUFFER  DATA  256F'0'       I/O BUFFER
KEY     TEXT   'KEY0001',LENGTH=7  RECORD KEY
```

LOAD - OPEN FILE FOR RECORD LOADING

The LOAD request builds an indexed access control block (IACB) associated with the file specified by the DSCB parameter. The address returned in the iacb variable is the address used to connect requests under this LOAD to this file.

To access the file by primary key, specify the primary file name as the DSCB parameter. On all subsequent requests, specify a primary key.

To access the file by secondary key, specify the secondary file name as the DSCB parameter. On all subsequent requests, specify a secondary key. The Indexed Access Method automatically opens the primary file when you specify a secondary file.

Note: The directory must be set up to reflect the relationship among the primary file and any secondary files.

LOAD opens the file for loading base records; the only acceptable processing requests in this mode are PUT, EXTRACT and DISCONN. Only one user of a file can use the LOAD function at one time.

If an error exit is specified, the error exit routine is executed whenever any Indexed Access Method request under this LOAD terminates with a positive return code.

Note: Return codes 1, 7, 8, and 22 are positive value return codes but they do not cause the error exit routine to be entered, even when ERREXIT is coded. Also the negative (warning) return codes do not cause error exits. For details on coding ERREXIT, see "LOAD - Open File for Record Loading," or "PROCESS - Open File" on page IG-122.

Syntax:

```
label      CALL      IAM,(LOAD),iacb,(dscb),(opentab),(mode)
```

```
Required:  iacb,dscb,opentab
```

```
Defaults:  mode=(SHARE)
```

Operands Description

iacb The label of a 1-word variable into which the address of the indexed access control block (IACB) is returned.

(dscb) The name of a valid DSCB. This name is DS n , where n is a number from 1 - 9, corresponding to a file defined by the PROGRAM statement. It can also be a name supplied by a DSCB statement. The CALL statement specifying LOAD causes the Indexed Access Method to open the index file in load mode.

(opentab) The label of a 3 word open table. The open table contains information used during this LOAD. The format of this table is as follows:

Offset	Field
0	SYSRTCD
2	ERREXIT
4	(0) reserved

Field Description

SYSRTCD A 1-word variable in which the return code from LOAD and from any system function (such as READ and WRITE) is placed when requested under this LOAD by the Indexed Access Method.

ERREXIT Your error exit routine address. If this address is zero, the error exit will not be taken. Note that error exits handle only positive return codes.

Note: Return codes 1, 7, 8, and 22 are positive return codes which do not cause the error exit routine to be entered, even if ERREXIT is coded.

RESERVED Must be 0 for LOAD requests.

(mode) Specifies shared or exclusive use of the file.

SHARE Allows shared read/write access by PROCESS requests.

ISHARE Allows shared read/write access by PROCESS requests with the independent processing flag on. No automatic update is performed on associated secondary indexes even if the auto-update flag is on in the directory for those secondary indexes.

For a secondary index, the index is opened as an independent file and the records returned are secondary index records, not user data records.

EXCLUSV You can access the file in exclusive mode (EXCLUSV) only if there are no outstanding PROCESS or LOAD requests. No other user can access the file while exclusive use is in effect.

IEXCLUSV You can access the file only if there are no outstanding PROCESS or LOAD requests. No other user can access the file while independent exclusive (IEXCLUSV) use is in effect. The I prefix on EXCLUSV mode prevents any automatic update functions on any associated secondary indexes, even if the auto-update

flag is on in the directory entry for those associated secondary indexes.

For a secondary index, the index is opened as an independent file and the records returned are secondary index records, not user data records.

LOAD Return Codes

Code	Condition
-79	Warning - File was opened and not closed during the last session. Normal processing continues
-75	Warning - File has either not been formatted, or the invalid indicator is on in the directory for that file
-57	Data set has been loaded
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
14	Invalid index block found - during processing an index block, an incorrect block type was found.
17	IAM is inactive - not enough storage available. Use \$IAMUT1 BF command to readjust storage size.
23	Insufficient number of IACBs, use BF command of \$IAMUT1 to allocate more
50	File opened exclusive
51	Data set already opened in load mode
52	File in use, cannot open exclusive
54	\$IAM buffer too small to process a file with this block size Use the BF command of \$IAMUT1 to increase the buffer size
55	Insufficient FCBs
56	Read error - FCB. Refer to system return code

Code	Condition
76	DSOPEN error occurred - The system error field in the open table contains the DSOPEN errors 21 - DSNAME,VOLUME not found 22 - VOLSERV error 23 - I/O error
77	Record save area not large enough - use \$IAMUT1 BF command to set maximum record size for secondary index processing
78	Attempted to open a secondary file for LOAD, file is not opened independently
230	Directory READ error from \$IAMDIR. Check system return code
234	Directory error - DSNAME,VOL not found in \$IAMDIR
243	Primary file failed to open on a secondary OPEN request

LOAD Example

The following example opens the file identified by 'DS3' for record loading in exclusive mode. The field named 'IACB' is set to the address of the IACB for this open. Subsequent requests use this field to refer to this file. The system return code is placed in the field named 'OPEN'. An error opening the file results in the routine named 'ERROR' being executed.

```

CALL    IAM,(LOAD),IACB,(DS3),(OPEN),(EXCLUSV)
.
.
.
IACB  DATA  F'0'
OPEN  DATA  F'0'          RETURN CODES
      DATA  A'ERROR'     ERROR EXIT ROUTINE ADDRESS
      DATA  F'0'          NOT USED

```

PROCESS - OPEN FILE

The PROCESS request builds an indexed access control block (IACB) associated with the file specified by the DSCB parameter. The address returned in the IACB variable is the address used to connect requests under this PROCESS to this file.

To access the file by primary key, specify the primary file name as the DSCB parameter. On all subsequent requests, specify a primary key.

To access the file by secondary key, specify the secondary file name as the DSCB parameter. On all subsequent requests, specify a secondary key. The Indexed Access Method automatically opens the primary file when you specify a secondary file.

Note: The directory must be set up to reflect the relationship between the primary file and any secondary files.

PROCESS opens the file for retrievals, updates, insertions, and deletions. Multiple users can PROCESS the same file. However, only one user at a time can use the LOAD function for a given file.

If ERREXIT is specified, the error exit routine is executed whenever any Indexed Access Method request under this PROCESS terminates with a positive return code.

Note: Return codes 1, 7, 8, and 22 are positive value return codes but they do not cause the error exit routine to be entered, even when ERREXIT is coded. Also the negative (warning) return codes do not cause error exits. For details on coding ERREXIT, see "LOAD - Open File for Record Loading" on page IG-118, or "PROCESS - Open File."

If EODEXIT is specified, the end-of-data exit routine is executed whenever a GETSEQ associated with PROCESS attempts to access a record after the last record in the file.

Syntax:

```
label      CALL  IAM,(PROCESS),iacb,(dscb),(opentab),(mode)

Required:  iacb,dscb,opentab
Defaults:  mode=(SHARE)
```

Operands Description

- iacb** The label of a 1-word variable into which the address of the indexed access control block (IACB) is returned.
- (dscb)** The name of a valid DSCB. This name is DS_n, where n is a number from 1 - 9, corresponding to a file defined by the PROGRAM

statement. It can also be a name supplied by a DSCB statement. The CALL statement specifying PROCESS, causes the Indexed Access Method to open the index file in process mode.

(opentab) The label of a 3 word open table. The open table contains information used during this PROCESS. The format of this table is as follows:

Offset	Field
0	SYSRTCD
2	ERREXIT
4	EODEXIT

Field Description

SYSRTCD A 1-word variable in which the return code from PROCESS and from any system function (such as READ and WRITE) is placed when requested under this PROCESS by the Indexed Access Method.

ERREXIT Your error exit routine address. If this address is 0, the error exit will not be used. Note that error exits handle only positive return codes.

EODEXIT Your end-of-data exit routine address. If this address is 0, the end-of-data exit will not be used.

(mode) Specifies shared or exclusive access to the file.

SHARE Allows shared read/write access by multiple PROCESS or LOAD requests.

ISHARE Allows shared read/write access by PROCESS requests with the independent processing flag on. No automatic update is performed on associated secondary indexes even if the auto-update flag is on in the directory for those secondary indexes.

For a secondary index, the index is opened as an independent file and the records returned are secondary index records, not user data records.

EXCLUSV The user can access the file only if there are no outstanding PROCESS or LOAD requests. No other user can access the file while EXCLUSV (exclusive access) is in effect.

IEXCLUSV You can access the file only if there are no outstanding PROCESS or LOAD requests. No other user can access the file while independent exclusive (IEXCLUSV) use is in effect. The I prefix on EXCLUSV

mode prevents any automatic update functions on any associated secondary indexes, even if the auto-update flag is on in the directory entry for those associated secondary indexes.

For a secondary index, the index is opened as an independent file and the records returned are secondary index records, not user data records.

PROCESS Return Codes

Code	Condition
-79	Warning - File was opened and not closed during the last session. Normal processing continues
-75	Warning - File has either not been formatted, or the invalid indicator is on in the directory for that file
-57	Data set has been loaded
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
17	IAM is inactive - not enough storage available. Use \$IAMUT1 BF command to readjust storage size.
23	Insufficient number of IACBs, use BF command of \$IAMUT1 to allocate more
50	File opened exclusive
51	Data set already opened in load mode
52	File in use, cannot open exclusive
54	\$IAM buffer too small to process a file with this block size Use the BF command of \$IAMUT1 to increase the buffer size
55	Insufficient FCBs
56	Read error - FCB. Refer to system return code
76	DSOPEN error occurred - The system error field in the open table contains the DSOPEN errors <ul style="list-style-type: none"> 21 - DSNAME, VOLUME not found 22 - VOLSERV error 23 - I/O error
77	Record save area not large enough - use \$IAMUT1 BF command to set maximum record size for secondary index processing
78	Attempted to open a secondary file for LOAD, file is not opened independently
230	Directory READ error from \$IAMDIR. Check system return code
234	Directory error - DSNAME, VOL not found in \$IAMDIR
243	Primary file failed to open on a secondary OPEN request

PROCESS Example

The following example opens the file identified by 'DS1' for general access in shared access mode. The field named 'IACB' is set to the address of the IACB for this open. Subsequent requests use this field to refer to this file. The system return code is placed in the field named 'OPENTAB'. An error opening the file results in the routine named 'ERROR' being executed. An end-of-data condition on a subsequent request results in the transfer of control to the code at the label 'END'.

CALL	IAM,(PROCESS),IACB,(DS1),(OPENTAB),(SHARE)	
.		
.		
.		
OPENTAB DATA	F'0'	RETURN CODES
DATA	A(ERROR)	ADDRESS OF ERROR EXIT ROUTINE
DATA	A(END)	ADDRESS OF EOD EXIT ROUTINE
IACB DATA	F'0'	

PUT - PUT RECORD INTO FILE

The PUT request processes the record that is in your buffer (buff) according to the way the file was opened (LOAD or PROCESS).

If the current open is for LOAD, the record must have a higher key than the highest key already in the file and only base record slots are used (refer to "Loading Base Records From An Application Program" on page IG-65 for a description of load mode). If the current open is for PROCESS, the record may have any key and is placed in key order in either a base record or in a free slot in the appropriate place in the file.

The PUTC request processes the record that is in your buffer (buff) according to the way the file was opened (LOAD or PROCESS) only if the record or file is not locked.

Syntax:

label	CALL	IAM, (PUT PUTC), iacb, (buff)
Required:	all	
Defaults:	none	

Operands Description

iacb	The label of a word containing the IACB address returned by PROCESS or LOAD.
(buff)	The label of the user area containing the record to be added to the file.

PUT Return Codes

Code	Condition
-90	Request cancelled because the request was conditional and a wait on a lock or buffer would be required
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
14	Invalid index block found - during processing an incorrect index block type was found, recreate the file
22	Address supplied by your program is not a valid IACB
60	Out of sequence or duplicate key (LOAD only)
61	End of file
62	Duplicate key found (PROCESS only)
70	No space for insert; reorganize the file
76	DSOPEN error occurred - The system error field in the OPEN table contains the DSOPEN errors: 21 - DSNAME,VOLUME not found 22 - VOLSERV error 23 - I/O error
90	Internal key save area temporarily in use by another request
100	Read error - check system return code
101	Write error - check system return code
230	Directory read error form \$IAMDIR. Check system return code
244	Error in opening auto-update on modification request
246	Auto-update processing an INSERT to a secondary failed, auto-update processing continues
248	I/O error on primary file during a secondary request

PUT Example

The following example puts the record in the area named 'BUFFER' into the file. The file is identified by the field named 'FILE1'.

```
CALL  IAM,(PUT),FILE1,(BUFFER)
      .
      .
      .
FILE1  DATA  F'0'           IACB ADDRESS RETURNED HERE
BUFFER DATA  256F'0'       I/O BUFFER
```

PUTDE - DELETE PREVIOUSLY READ RECORD

The PUTDE request deletes a record from an indexed file. The record must have been previously retrieved by a GET or GETSEQ in update mode. Deleting the record creates free space in the file. The PUTDE releases the lock placed on the record by the GET or GETSEQ.

The PUTDEC request deletes a record from an indexed file only if the block or buffer is not locked.

Syntax:

label	CALL	IAM,(PUTDE PUTDEC),iacb,(buff)
Required:	all	
Defaults:	none	

Operands Description

- iacb** The label of a word containing the IACB address returned by PROCESS.
- (buff)** The name of the area containing the record previously retrieved by GET or GETSEQ.

PUTDE Return Codes

Code	Condition
-90	Request cancelled because the request was conditional and a wait on a lock or buffer would be required
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
14	Invalid index block found - during processing an incorrect index block was found. Recreate the file
22	Address supplied by your program is not a valid IACB
76	DSOPEN error occurred - The system error field in the OPEN table contains the DSOPEN errors: 21 - DSNAME, VOLUME not found 22 - VOLSERV error 23 - I/O error
85	Key was modified by user
100	Read error - check system return code
101	Write error - check system return code
230	Directory read error from \$IAMDIR. Check system return code
242	Secondary index is out of sync with primary file. Must rebuild file to get back in sync.
244	Error in opening auto-update on modification request
245	Auto update PUTDE to a secondary failed, Auto-update processing continues.
247	During auto-update processing a GETSEQ to a secondary failed, auto-update processing continues.
248	I/O error on primary file during a secondary request.

PUTDE Example

The following example deletes the record in the area named 'BUFFER' from the file. The record was read with either a GET or GETSEQ request in update mode. The file is identified by the field named 'FILE1'.

```
CALL  IAM,(PUTDE),FILE1,(BUFFER)
.
.
.
FILE1  DATA  F'0'          IACB ADDRESS FROM PROCESS
BUFFER DATA  256F'0'      I/O BUFFER
```

PUTUP - UPDATE RECORD

The PUTUP request replaces the record in the file with the record in your buffer. The record must have been retrieved by a GET or GETSEQ in update mode. You must not change the primary key field in the record or the contents of the key area in your program returned by the GET or GETSEQ request. The Indexed Access Method checks for and prohibits primary key modification. The PUTUP releases the lock placed on the record by the GET or GETSEQ.

The PUTUPC request replaces the record in the file with the record in your buffer only if the record, block, or buffer is not locked.

Syntax:

```
label      CALL      IAM,(PUTUP|PUTUPC),iacb,(buff)
```

```
Required:  all
```

```
Defaults:  none
```

Operands Description

iacb The label of a word containing the IACB address returned by PROCESS.

(buff) The label of the user area containing the record to replace the one previously retrieved.

PUTUP Return Codes

Code	Condition
-90	Request cancelled because the request was conditional and a wait on a lock or buffer would be required
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
14	Invalid index block found - during processing an incorrect index block was found. Recreate the file
22	Address supplied by your program is not a valid IACB
76	DSOPEN error occurred - The system error field in the OPEN table contains the DSOEEN errors: 21 - DSNAME,VOLUME not found 22 - VOLSERV error 23 - I/O error
85	Key was modified by user
100	Read error - check system return code
101	Write error - check system return code
230	Directory read error form \$IAMDIR. Check system return code
242	Secondary index is out of sync with primary file. Must rebuild file to get back in sync.
244	Error in opening auto-update on modification request
245	Auto update PUTDE to a secondary failed, auto-update processing continues.
246	Auto-update processing an INSERT to a secondary failed, auto-update processing continues
247	During auto-update processing a GETSEQ to a secondary failed, auto-update processing continues.
248	I/O error on primary file during a secondary request.

PUTUP Example

The following example puts the updated record in the area named 'BUFFER' back into the file. The record was read with either a GET or GETSEQ request in update mode. The file is identified by the field named 'FILE1'.

```
CALL  IAM,(PUTUP),FILE1,(BUFFER)
      .
      .
      .
FILE1  DATA  F'0'           IACB ADDRESS FROM PROCESS
BUFFER DATA  256F'0'       I/O BUFFER
```

RELEASE - RELEASE RECORD

The RELEASE request frees a record that has been locked by a GET or GETSEQ for update. A record lock is normally released by a PUTUP or PUTDE. The RELEASE request is useful for freeing the locked record when the update need not be completed. RELEASE is valid only when a record is locked for update.

Syntax:

```
label      CALL      IAM,(RELEASE),iacb
```

```
Required:  all  
Defaults:  none
```

Operands Description

iacb The label of a word containing the IACB address returned by PROCESS.

RELEASE Return Codes

Code	Condition
-1	Successful
7	Link module in use, synchronize use of link module with the program
8	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
22	Address supplied by your program is not a valid IACB

RELEASE Example

The following example releases the record that was read with either a GET or GETSEQ request in update mode. The file is identified by the field named 'FILE1'.

```
CALL IAM,(RELEASE),FILE1
.
.
FILE1 DATA F'0' IACB ADDRESS FROM PROCESS
```

EDL CALL FUNCTIONS SYNTAX SUMMARY

Following is a summary of the syntax of the EDL CALL instructions used to invoke the functions provided by the Indexed Access Method.

```
label CALL IAM,(DELETE|DELETC),iacb,(key)
label CALL IAM,(DISCONN),iacb
label CALL IAM,(ENDSEQ),iacb
label CALL IAM,(EXTRACT),iacb,(buff),(size),(type)
label CALL IAM,(GET|GETC|GETR|GETCR),iacb,(buff),(key),(mode/krel)
label CALL IAM,(GETSEQ|GETSEQC|GETSEQCR|GETSEQR),iacb,(buff),
      (key),(mode/krel)
label CALL IAM,(LOAD),iacb,(dscb),(opentab),(mode)
label CALL IAM,(PROCESS),iacb,(dscb),(opentab),(mode)
label CALL IAM,(PUT|PUTC),iacb,(buff)
label CALL IAM,(PUTDE|PUTDEC),iacb,(buff)
label CALL IAM,(PUTUP|PUTUPC),iacb,(buff)
label CALL IAM,(RELEASE),iacb
```

INDEXED ACCESS METHOD RETURN CODES SUMMARY

Return Code	Condition
-90	Request cancelled because the request was conditional and a wait on a lock or buffer would be required
-85	Record to be deleted not found
-80	End of data
-79	Warning - File was opened and not closed during the last session, normal processing continues
-75	Warning - File has either not been formatted or the invalid indicator is on in the directory for that file
-58	Record not found
-57	Data set has been loaded
-1	Successful completion
01	Invalid function specified on CALL to \$IAM
07	Link module in use, synchronize use of link module with the program
08	Load error for \$IAM, verify \$IAM exists and enough storage is available to load it
10	Invalid request
12	Data set shut down due to error; see version 2 guide, error recovery
13	A required module is not included in \$IAM
14	Invalid index block found - during processing an incorrect index block type was found, recreate the file
17	IAM is inactive - not enough storage available Use \$IAMUT1 BF command to readjust storage size
22	Address supplied by your program is not a valid IACB
23	Insufficient number of IACBs, use BF command of \$IAMUT1 to allocate more
50	Data set is opened for exclusive use, cannot be opened exclusively
51	Data set already opened in load mode
52	Data set is opened, cannot be opened exclusively
54	\$IAM buffer too small to process a file with this block size Use the BF command of \$IAMUT1 to increase the buffer size
55	Get storage error - FCB
56	READ error - FCB, refer to system return code
60	Out of sequence or duplicate key
61	End of file
62	Duplicate key found in process mode
65	Inconsistent free space parameters RSVIX/RSVBLK or FPOOL/DYN were specified without one or the other
70	No space for insert. Reorganize the file

Return Code	Condition
76	DSOPEN error occurred - The system error field in the OPEN table contains the DSOPEN errors: 21 - DSNAME,VOLUME not found 22 - VOLSERV error 23 - I/O error
77	Record save area not large enough - use \$IAMUT1 BF command to set maximum record size for secondary file processing
78	Attempted to open a secondary file for LOAD, file is not opened independently
80	FCB WRITE error during DELETE processing - see system return code
85	Key field modified by user
90	Internal key save area temporarily in use by another request
100	READ error - check system return code
101	WRITE error - check system return code
110	WRITE error - data set closed
120	Invalid EXTRACT type
122	File does not contain FCB extension
123	Cannot extract paging statistics. Data paging is not active
150	Not enough storage available for data paging
200	Error occurred while accessing the primary file
230	Directory read error form \$IAMDIR
231	\$IAMQCB not found. Check sysgen for include of \$IAMQCB
234	Directory error - DSNAME,VOL not found in \$IAMDIR
242	Secondary index is out of sync with primary file. Must rebuild file to get back in sync.
243	Primary file failed to open on secondary open request
244	Primary I/O error
245	Error in opening an auto-update modification request
246	Auto-update processing an INSERT to a secondary failed, auto-update processing continues
247	During auto-update processing a GETSEQ to a secondary failed, auto-update processing continues
248	I/O error on primary file during a secondary request
249	GET UPDATE error occurred trying to update a bad RBN

Note: For return codes 243 through 249, multiple errors may have occurred. Use \$ILOG to display the errors.

CHAPTER 8. THE \$IAMUT1 UTILITY

This chapter describes how to use the \$IAMUT1 utility to build and maintain your indexed files. Each command is described, including its function, parameters, and an example of how to use it. The file definition parameters are also described.

The chapter is arranged in alphabetic order. Following is the list of commands and the location of their descriptions:

- "BF - Tailor the Indexed Access Method Buffers" on page IG-144
- "DF - Define Indexed File" on page IG-146
- "DI - Display Parameter Values" on page IG-149
- "DR - Invoke Secondary Index Directory Functions" on page IG-150
- "EC - Control Echo Mode" on page IG-160
- "EF - Display Existing Indexed File Characteristics" on page IG-161
- "LO - Load Indexed File" on page IG-162
- "NP - Deactivate Paging" on page IG-166
- "PG - Select Paging" on page IG-167
- "PP - Define Paging Partitions" on page IG-168
- "PS - Get Paging Statistics" on page IG-169
- "RE - Reset Parameters" on page IG-170
- "RO - Reorganize Indexed File" on page IG-171
- "SE - Set Parameters" on page IG-173
- "UN - Unload Indexed File" on page IG-184

The subcommands of the directory function (DR), are listed alphabetically under the DR description. Those subcommands are:

- AL - allocate/reallocate directory
- EN - end directory function
- DE - delete entry
- IE - insert entry
- LE - list entries
- UE - update entry

\$IAMUT1

\$IAMUT1 can be invoked using the \$L command, \$JOBUTIL, or the Session Manager. \$IAMUT1 functions use dynamic storage for work and buffer areas. The \$IAMUT1 utility is shipped with sufficient dynamic storage to handle input and output block sizes of up to 512 bytes. This enables you to define an indexed file with a maximum block size of 512 bytes, and to load, unload, and reorganize indexed files with a maximum block size of 512 bytes. \$IAMUT1 determines if enough dynamic storage has been provided. If sufficient storage has not been provided, \$IAMUT1 displays a message. In order to handle large blocks of data, a larger dynamic storage area will have to be provided to \$IAMUT1. Additional dynamic storage can be provided by one of two ways: provide the storage parameter on the \$L command, or use the SS command of the \$DISKUT2 utility.

The load, unload and reorganize functions use the entire dynamic storage available to minimize the number of disk I/O operations. Improved performance, therefore, can be obtained by specifying as large a dynamic area as possible.

\$IAMUT1 COMMANDS

The commands available under \$IAMUT1 are listed below. To display this list at your terminal, enter a question mark in response to the prompting message ENTER COMMAND (?):.

The command descriptions in this chapter are arranged in alphabetic order.

ENTER COMMAND (?): ?

EC - SET/RESET ECHO MODE
EF - DISPLAY EXISTING FILE CHARACTERISTICS
DR - SECONDARY INDEX DIRECTORY FUNCTIONS
EN - END THE PROGRAM

SE - SET DEFINE PARAMETERS
DF - DEFINE AN INDEXED FILE
DI - DISPLAY CURRENT SE PARAMETERS
RE - RESET CURRENT VALUES FOR DEFINE

LO - LOAD INDEXED FILE FROM SEQUENTIAL FILE
RO - REORGANIZE INDEXED FILE
UN - UNLOAD INDEXED FILE TO SEQUENTIAL FILE

PG - SELECT DATA PAGING
NP - DESELECT DATA PAGING
PP - DEFINE PAGING PARTITIONS
PS - DATA PAGING STATISTICS
BF - SET BUFFER SIZES

ENTER COMMAND (?):

After the commands are displayed, you are again prompted with ENTER COMMAND (?):. Respond with the command you wish to use.

BF

BF - TAILOR THE INDEXED ACCESS METHOD BUFFERS

The BF command specifies the amount of storage that the Indexed Access Method (\$IAM) is to use for buffers and control blocks and the maximum record size for any file with a secondary index.

BF prompts you for each of the following parameters by displaying the current value and accepting new settings.

BUFFER SIZE Indicates the amount of storage (in bytes) to be used for the central buffer. Use the following formula to calculate your minimum buffer size:

Buffer Size = (2 x blocksize) + (28 x blocksize/256)
where: blocksize = maximum block size

NUMBER OF IACBs Indicates the number of the IACBs. The maximum number of IACBs is 64. There is an IACB associated with each PROCESS or LOAD that is issued. When calculating the number of IACBs you should consider the number of concurrent users you may have at any one time.

NUMBER OF FCBS Indicates the number of FCBS. The maximum number of FCBS is 64. There is one FCB for every file that is open. When calculating the number of FCBS you should consider the maximum number files that might be open at a given time.

MAXIMUM RECORD SIZE Indicates the maximum record size of any file with an associated secondary index. If no files have a secondary index, this value can be zero. The actual amount of storage reserved as a result of this parameter is twice the value specified plus 8 bytes.

None of these take effect until the next time the Indexed Access Method is loaded.

BF Command Example

This example sets the central buffer size to 540 bytes, leaves the number of IACBs at 3, leaves the number of FCBS at 3, and sets the maximum record size of any file with a secondary index to 120 bytes.

```
ENTER COMMAND (?): BF
PARAMETER          DEFAULT    NEW VALUE
BUFFER SIZE        1080      : 540
NUMBER OF IACBs    3         :
NUMBER OF FCBS     3         :
MAXIMUM RECORD SIZE 256      : 120
VALUE(S) SET
STORAGE FOR $IAM HAS BEEN SET TO 2048
BECOMES EFFECTIVE ON NEXT LOAD OF $IAM

ENTER COMMAND (?):
```

DF

DF - DEFINE INDEXED FILE

The DF command allocates, defines, and formats an indexed file. The DF function will optionally invoke the load or reorganize function for you. Before entering DF, you must use the SE command to set up parameters that determine the size and format of the indexed file. The DF command uses those SE parameters to optionally allocate and format the file. The DF function can be invoked at the end of the SE function.

The allocate step consists of using the file size computed during the SE step to dynamically allocate the file. If the file already exists, the size is verified to ensure that it is large enough. The define step consists of writing the file control block (FCB) and its extension to the indexed file. Finally, the optional format step initializes all records in the indexed file to provide an empty structured file.

INVOKING THE LOAD AND REORGANIZE FUNCTIONS FROM DF: You can invoke the LOAD or REORGANIZE functions directly from the DF (or SE) command. If you invoke these functions, DF does not format the file because LOAD and REORGANIZE will format the file. If you do not invoke the LOAD or REORGANIZE function, DF formats the file so you can load the file using an application program or the LO command.

Notes:

1. You can use the LOAD/REORGANIZE command later to load the file, if you do not invoke it from the DF command.
2. An application program cannot access an unformatted indexed file.
3. The prompt for the load/reorganize function occurs before the file is actually defined.
4. A secondary index file cannot be loaded with the LO command, though it can be reorganized using the RO function.

Defining the File

The define function prompt for the file to be allocated. If the file already exists, its size is checked. If the size is at least as large as needed, DF prompts you as to whether the file should be reused as follows:

DF

```
ENTER COMMAND (?): DF
ENTER DATA SET (NAME,VOLUME) : IAMFILE,EDX003
DATA SET ALREADY EXISTS AND IS LARGE ENOUGH
DO YOU WISH TO REUSE IT (Y/N)? : Y
```

If the file exists, but it is not as large as needed, you have the option of deleting and reallocating it as shown in the following example:

```
ENTER COMMAND (?): DF
ENTER DATA SET (NAME,VOLUME) : MASTER,VOL123
DATA SET ALREADY EXISTS AND IS TOO SMALL
DELETE AND REALLOCATE (Y,N)? : Y
DELETE AND REALLOCATE COMPLETED
```

If the file does not exist, it is allocated as follows:

```
ENTER COMMAND (?): DF
ENTER DATA SET (NAME,VOLUME) : MASTER,VOL123
NEW DATA SET IS ALLOCATED
```

Using Immediate Write-Back

DF prompts you to select whether or not you want to use the immediate write-back option. Immediate write-back has the same effect on primary or secondary indexed files.

Each request to insert, delete, or update a data record, causes the affected blocks to be read into the Indexed Access Method buffer. The actual modification to the block is performed in the buffer.

If you enter N to the immediate write-back prompt, file modifications are held in the main storage buffer and not written back to the indexed file until the buffer space is needed for another block or until the file is closed. If the device where the file resides was powered off before the block was written back to the file, the modification to the file would not have been performed.

DF

If you enter Y to the immediate write-back prompt, you are assured that the changed block is written back to the file immediately.

The prompt is as follows:

DO YOU WANT IMMEDIATE WRITE-BACK? Y

DF Command Example

The following example shows a use of the DF command to define a file named MASTER on volume VOL123. Immediate write-back is selected and the request to invoke LOAD or REORGANIZE is indicated.

```
ENTER COMMAND (?): DF
ENTER DATA SET (NAME,VOLUME) : MASTER,VOL123
NEW DATA SET IS ALLOCATED
DO YOU WANT IMMEDIATE WRITE-BACK? Y
INVOKE LOAD(L), REORGANIZE(R) OR END(E) AFTER CURRENT FUNCTION? L
DEFINE IN PROGRESS
DATA SET SIZE IN EDX RECORDS:          17
INDEXED ACCESS METHOD RETURN CODE:      -1
SYSTEM RETURN CODE:                    -1
PROCEED WITH LOAD/REORGANIZE (Y/N)
```

DI - DISPLAY PARAMETER VALUES

DI displays the current parameter values entered during the current session of \$IAMUT1 SE command. The parameter values can be used to format a file using the DF command or they can be modified by reusing the SE command.

Note: You can also use the EF command to display the parameters of an existing file.

The following example shows a use of the DI command.

```
ENTER COMMAND (?): DI
CURRENT VALUES FOR SE COMMAND ARE:
FILE TYPE = PRIMARY
BASEREC          100
BLKSIZE          256
RECSIZE          80
KEYSIZE          28
KEYPOS           1
FREEREC          1
FREEBLK          10
RSVBLK           NULL
RSVIX            0
FPOOL            NULL
DELTHR           NULL
DYN              NULL
```

For a secondary file, the record size is not displayed.

DR

DR - INVOKE SECONDARY INDEX DIRECTORY FUNCTIONS

The DR command provides access to secondary index directory functions. Those functions are made available by replying DR when \$IAMUT1 requests "ENTER COMMAND (?):". You can then respond to the "ENTER DIRECTORY COMMAND (?):" with a subcommand. To obtain a list of the available subcommands, reply with a question mark (?) as follows:

```
ENTER COMMAND (?): DR

ENTER DIRECTORY COMMAND (?): ?

AL - ALLOCATE/REALLOCATE DIRECTORY
LE - LIST ENTRIES
ID - INSERT ENTRY
DE - DELETE ENTRY
UE - UPDATE ENTRY
EN - END DIRECTORY FUNCTION

ENTER DIRECTORY COMMAND (?):
```

The directory function commands are arranged in alphabetic order as DR - xx, where xx is the two letter directory subcommand. Their specific locations are listed below:

- "AL - Allocate Directory" on page IG-151
- "DE - Delete Directory Entry" on page IG-153
- "EN - End Directory Function" on page IG-154
- "IE - Insert Entry" on page IG-155
- "LE - List Entries" on page IG-156
- "UE - Update Directory Entry" on page IG-158

AL - ALLOCATE DIRECTORY

The AL subcommand allocates a directory for secondary indexes. If a directory already exists, this subcommand gives the option to delete and reallocate it.

Note: To use this subcommand, you must first use the DR command.

You are prompted to enter the maximum number of directory entries. Enter the number of entries you want the directory to be able to hold. Each entry describes a primary file or secondary index. The maximum number of entries defaults to 47.

The directory, \$IAMDIR, is always allocated on the IPL volume.

The following example shows a use of the AL subcommand to allocate a new directory with a capacity of 10 entries:

```
ENTER DIRECTORY COMMAND (?): AL
MAX # OF DIRECTORY ENTRIES: 10
THE DIRECTORY DATA SET REQUIRES 1 EDX RECORDS, CONTINUE (Y/N/EN)? Y
DIRECTORY DATA SET ALLOCATED: $IAMDIR,EDX002
```

The next example assumes a directory already exists and allocates a new one.

```
ENTER DIRECTORY COMMAND (?): AL
DIRECTORY EXISTS, OPTIONS ARE:
BN - BUILD NEW DIRECTORY
AS - ADJUST SIZE
EN - END DIRECTORY ALLOCATE

ENTER OPTION: BN

ALL DIRECTORY ENTRIES WILL BE DELETED, CONTINUE (Y/N)? Y
MAX # OF DIRECTORY ENTRIES: 20
THE DIRECTORY DS REQUIRES      2 EDX RECORDS, CONTINUE (Y/N/EN) ? Y

DIRECTORY DATA SET ALLOCATED: $IAMDIR,EDX002
```

The following example, adjusts the size of the directory data set. All existing entries will be retained.

DR - AL

ENTER DIRECTORY COMMAND (?): AL

DIRECTORY EXISTS, OPTIONS ARE:

BN - BUILD NEW DIRECTORY

AS - ADJUST SIZE

EN - END DIRECTORY ALLOCATE

ENTER OPTION: AS

MAX # OF DIRECTORY ENTRIES: 1

THE DIRECTORY DS REQUIRES 1 EDX RECORDS, CONTINUE (Y/N/EN) ? Y

DIRECTORY DATA SET ALLOCATED: \$IAMDIR,EDX002

DE - DELETE DIRECTORY ENTRY

The DE subcommand deletes an entry from the directory. If you delete a primary entry, all associated secondary index entries are also deleted.

Note: To use this subcommand, you must first use the DR command.

The following example shows the deletion of the directory entry for the file named MASTER on the volume named VOL123. MASTER is a primary index file entry which has secondary indexes associated with it.

```
ENTER DIRECTORY COMMAND (?): DE
ENTRY (DSNAME,VOLUME): MASTER,VOL123
ASSOCIATED SECONDARY ENTRIES WILL BE DELETED, CONTINUE (Y/N)? Y
DELETE SUCCESSFUL, NUMBER OF ENTRIES DELETED: 2
```

The following example shows the deletion of the directory entry for a file named MASTER, on the volume named VOL123. MASTER is a primary index file entry which no longer has any secondary indexes associated with it.

```
ENTER DIRECTORY COMMAND (?): DE
ENTRY (DSNAME,VOLUME): MASTER
ENTRY FOR MASTER ,EDX002 WILL BE DELETED, CONTINUE (Y/N)? Y
DELETE SUCCESSFUL, NUMBER OF ENTRIES DELETED: 1
```

DR - EN

EN - END DIRECTORY FUNCTION

The EN subcommand terminates the directory functions (DR) and returns to \$IAMUT1 for your next command.

IE - INSERT ENTRY

The IE subcommand inserts a new entry into the secondary index directory. It is used to insert either a primary or secondary entry. However, the primary entry must be inserted before any of its secondary entries can be inserted.

For a primary entry, enter the data set name and volume of the file for which the entry is being inserted. Specify N when asked "IS THIS A SECONDARY ENTRY (Y/N)?."

For secondary entries, enter the data set name and volume of the secondary index for which the entry is being inserted and specify that it is a secondary index. You are then prompted for additional information.

Specify the name of the primary index file which the secondary index is to be associated with. You can select automatic update, which indicates that any change to a primary file is to be reflected in the secondary index. The default for automatic update is yes.

The following example inserts a directory entry for a primary index file:

```
ENTER DIRECTORY COMMAND (?): IE
ENTRY (DSNAME,VOLUME): TOMPRI,EDX002
IS THIS A SECONDARY ENTRY (Y/N)? N

DIRECTORY INSERT SUCCESSFUL
```

The following example inserts a directory entry for a secondary index named 'TOMSEC1,EDX002' which is to be associated with the primary index file 'TOMPRI,EDX002'. Automatic update is selected.

```
ENTER DIRECTORY COMMAND (?): IE
ENTRY (DSNAME,VOLUME) TOMSEC1,EDX002
IS THIS A SECONDARY ENTRY? Y

ASSOCIATED PRIMARY ENTRY (DSNAME,VOLUME): TOMPRI,EDX002
AUTO-UPDATE (Y/N)? Y
```

Note: To use this subcommand, you must first use the DR command.

DR - LE

LE - LIST ENTRIES

The LE subcommand lists the contents of one or more directory entries. Specify the name of a primary indexed file to get information about that file and its secondary indexes. Specify the name of a secondary index to get information about only that secondary index. To obtain a complete list of all information in the directory, just press the Enter Key without supplying any data set name or volume.

Note: To use this subcommand, you must first use the DR command.

The following example lists the directory entries related to the primary file named 'TOMPRI' on volume 'EDX002'.

```
ENTER DIRECTORY COMMAND (?): LE
ENTRY (DSMANE,VOLUME) BLANK=ALL: TOMPRI

      PRIMARY INDE-          AUTO
DSNAME  VOLUME DATA SET PENDENT  INVALID  UPDATE
-----
TOMPRI  EDX002 YES      NO      ****    ****
TOMSEC1 EDX002 NO       NO      YES     YES
TOMSEC2 EDX002 NO       NO      YES     NO

NUMBER OF DIRECTORY ENTRIES USED =      6
NUMBER OF AVAILABLE ENTRY SLOTS =     41
DIRECTORY LIST COMPLETED
```

The following example lists all directory entries.

```
ENTER DIRECTORY COMMAND (?): LE
ENTRY (DSNAME,VOLUME) BLANK=ALL:

      PRIMARYU  INDE-
DSNAME  VOLUME DATA SET  PENDENT  INVALID  AUTO
                                                UPDATE

EDXIAM  EDX003 YES      NO      ****   ****
EDXIAMS1 EDX003 NO       NO      YES    YES

TOMPRI  EDX002 YES      NO      ****   ****
TOMSEC1 EDX002 NO       NO      YES    YES
TOMSEC2 EDX002 NO       NO      YES    NO

NUMBER OF DIRECTORY ENTRIES USED =      6
NUMBER OF AVAILABLE ENTRY SLOTS =     41
DIRECTORY LIST COMPLETED
```

DR - UE

UE - UPDATE DIRECTORY ENTRY

The UE subcommand updates an entry in the secondary index directory. You can use this command as follows:

- Specify null values for parameters to remain unchanged (press the Enter key when you are prompted for them).
- Enter new values for parameters to be modified.

Note: You cannot change a primary entry to a secondary entry or a secondary entry to a primary entry. To do this, you must delete the old entry and insert a new one.

The following example updates a primary directory entry named 'MASTER,VOL123', changes the volume name from VOL123 to EDX002 and leaves the DSNAME MASTER as it is.

```
ENTER DIRECTORY COMMAND (?): UE
ENTRY (DSNAME,VOLUME) MASTER,VOL123
THIS IS A PRIMARY ENTRY
IN THE FOLLOWING, ENTER NEW VALUE OR,
ENTER NULL LINE TO RETAIN (PRESENT VALUE)
DSNAME (MASTER):
VOLUME (VOL123): EDX002
INDEPENDENT (N):
DIRECTORY UPDATE SUCCESSFUL
```

The following example updates a secondary directory entry named 'MASTER,VOL123', changes the VOLUME name to EDX002 and leaves the DSNAME MASTER as it is. It sets automatic update, leaves the independent processing flag as it is, and sets the invalid indicator off.

```
ENTER DIRECTORY COMMAND (?): UE

ENTRY (DSNAME,VOLUME) MASTER,VOL123

THIS IS A SECONDARY ENTRY
IN THE FOLLOWING, ENTER NEW VALUE OR,
ENTER NULL LINE TO RETAIN (PRESENT VALUE)

DSNAME (MASTER):
VOLUME (EDX123): EDX002
INDEPENDENT (N):
INVALID INDICATOR (Y): N
AUTO-UPDATE (Y): Y

DIRECTORY UPDATE ENDED
```

Note: To use this subcommand, you must first use the DR command.

EC

EC - CONTROL ECHO MODE

EC enables you to enter or leave echo mode. When in echo mode, all \$IAMUT1 input and output is logged on the \$SYSPRTR device. This enables you to save information about the files you maintain using \$IAMUT1. When in echo mode, all input and output is logged until either the current utility session is ended or echo mode is reset by use of the EC command. Echo mode is off when \$IAMUT1 is loaded.

Note: Input and output from \$DISKUT3 is not logged.

The following examples show the commands to set and reset echo mode:

```
ENTER COMMAND (?): EC
DO YOU WANT ECHO MODE? (Y/N)?: Y   (Set echo mode)
FUNCTION COMPLETED

ENTER COMMAND (?): EC
DO YOU WANT ECHO MODE? (Y/N)?: N   (Reset echo mode)
FUNCTION COMPLETED
```

EF - DISPLAY EXISTING INDEXED FILE CHARACTERISTICS

The EF command displays the file definition parameters that were used to set up the file. The information is obtained from the FCB Extension block.

EF Command Example

This example shows how to display the file parameters used to set up the file.

```
ENTER COMMAND (?): EF
EXHIBIT FUNCTION ACTIVE
ENTER DATASET (NAME,VOLUME): EDXIAM1,EDX003

FILE TYPE = PRIMARY
BASEREC          20
BLKSIZE          256
RECSIZE          80
KEYSIZE          4
KEYPOS           1
FREEREC          0
FREEBLK          0
RSVBLK           NULL
RSVIX            0
FPOOL            NULL
DELTHR           NULL
DYN              10
EXHIBIT FUNCTION COMPLETED
```

LO

LO - LOAD INDEXED FILE

LO loads a primary indexed file from a sequential (blocked or unblocked) input file. (A secondary indexed file must be loaded by using the DF or SE command). A primary indexed file can be loaded in one of two environments. Loading an empty file is referred to as the initial load. For an indexed file that already contains some records, the LO command can be used to add records with higher keys (keys of higher value than those already in the indexed file). This is called load in **extend** environment.

Blocks are read from the sequential file with the EDL READ instruction and de-blocking is performed, if necessary. In the initial load environment, data records are formatted into Indexed Access Method blocks and written to the indexed file with the EDL WRITE instruction. Corresponding index blocks are written as required. The remainder of the indexed file is formatted if formatting was not completed during the DF function. In the extend environment, records are loaded into the indexed file using Indexed Access Method PUT requests.

The sequential input file can contain blocked or unblocked records. For a description of blocked and unblocked sequential data sets, see "Blocked and Unblocked Sequential Data Sets" on page IG-164. The records in the sequential file must be in ascending order by the data contained in the key field. If a record with a duplicate or out of sequence key is found, you are given the option to either omit the record and continue loading, or to end loading. The indexed file must have been defined by using the SE and DF commands before using the LO command.

Your response to the prompt message "ENTER INPUT BLOCKSIZE", defines to the LO command whether the input is a blocked or unblocked sequential file. A null response to the prompt "ENTER INPUT BLOCKSIZE" indicates an unblocked input file and the block size is then calculated using the input record size value, rounded up to the next 256-byte multiple value. If the actual block size value is entered as your response to this prompt, a blocked sequential input file is indicated.

The record lengths of the input and output files do not have to be the same. When the indexed file is opened, the record length is displayed on the terminal. At this point, you can specify the record length of the sequential file if it is different than that of the indexed file. If the indexed file records are longer than the sequential file records, the loaded records are left justified and filled with binary zeroes. If the indexed file records are shorter than the sequential file records, the following message appears on the terminal:

LO

INPUT REC GT OUTPUT REC. TRUNCATION WILL OCCUR.
OK TO PROCEED?

Reply 'Y' to proceed (records will be truncated).

Reply 'N' to terminate the load function.

If the end of the input sequential file is reached, you can continue loading from another sequential file. You are asked if there is more data to load. If you reply yes (Y), you are prompted for the file and volume name of the new input sequential file to use. The load operation continues, putting the first record of the new input sequential file in the next available record slot of the indexed file.

Note: The record lengths and block sizes of subsequent input files are assumed to be the same as the initial input file.

If the end of input file is reached and you do not name another input file, the load operation is complete.

Note: If you are loading the indexed file from a tape file, \$IAMUT1 does not close the tape file upon completion of the load. Use the \$VARYOFF command to close the tape file (refer to the Operator's Reference for a description of the \$VARYOFF command).

The following example shows use of the LO command:

```
ENTER COMMAND (?): LO

LOAD ACTIVE
ENTER OUTPUT DATASET (NAME,VOLUME): IAMFILE,EDX003
$FSEDIT FILE RECSIZE = 128
INPUT RECORD ASSUMED TO BE      80 BYTES. OK?: Y
ENTER INPUT BLOCKSIZE (NULL = UNBLOCKED):
ENTER INPUT DATASET (NAME,VOLUME): SEQ01,EDX003
LOAD IN PROCESS

END OF INPUT DATASET
ANY MORE DATA TO BE LOADED?: N
      6 RECORDS LOADED
LOAD SUCCESSFUL
```



LO

Blocked and Unblocked Sequential Data Sets

The LO (load) function of \$IAMUT1 will accept either blocked or unblocked sequential data sets as input when loading an indexed file. The UN (unload) function will either block or unblock data as requested when unloading an indexed file to a sequential data set.


UNBLOCKED SEQUENTIAL DATA SET: An unblocked sequential data set contains one record in each block. The blocksize must be a multiple of 256 bytes. The record size must be equal or less than the block size. A block can span one or more EDX records.

The following diagram illustrates the relationship of a data record of 300 bytes to a block size of 512 bytes in an unblocked data set.



BLOCKED SEQUENTIAL DATA SET: In a blocked sequential data set a block can contain multiple logical records. The block size must be a multiple of 256 bytes. The record size must be equal to or less than the block size. A block can span one or more EDX record.

The following diagram illustrates 6 data records of 80 bytes each within a block of 512 bytes in a blocked data set.



Both the blocked and unblocked forms of sequential data sets, used by the utility, are compatible with the language processors, Sort/Merge and data

sets produced by \$FSEDIT. If you use the EDX edit utilities to prepare your data records for input, remember that these utilities put one 80-byte line from \$FSEDIT into a 128-byte \$FSEDIT record. Two of these 128-byte records are then used to form one 256-byte EDX record. When you use such a data set as sequential input for the LO (load) function, specify the record length as 128 and the block size as 256. If your indexed file is defined as having a record length of 80, you will receive the message "TRUNCATION WILL OCCUR." This is acceptable because Indexed Access Method strips off the extra bytes added by \$FSEDIT.

The last block of a blocked sequential data set may not have enough records for a full block. In this case, all of the unused space in the block is set to binary zeroes.

Invoking the LOAD and REORGANIZE Functions

You can invoke the LOAD or REORGANIZE functions directly from the DF command. If you invoke these functions, DF does not format the file because LOAD and REORGANIZE will do it. If you do not invoke the LOAD or REORGANIZE function, DF formats the file so you can load the file using an application program or \$IAMUT1 at a later time.

Notes:

1. You can use the LOAD/REORGANIZE command later to load the file, if you do not invoke it from the DF command.
2. An application program cannot access an unformatted indexed file.
3. The prompt for the load/reorganize function occurs before the define step.

NP

NP - DEACTIVATE PAGING

The NP command directs that data paging be deselected the next time the Indexed Access Method is loaded.

Page area sizes are not affected by this command.

NP Command Example

This example shows how to indicate data paging is to be deselected on the next invocation of the Indexed Access Method.

```
ENTER COMMAND (?): NP
DATA PAGING MARKED AS NOT ACTIVE
BECOMES EFFECTIVE ON NEXT LOAD OF $IAM
```

PG - SELECT PAGING

The PG command directs that data paging be selected the next time the Indexed Access Method is loaded.

Page area sizes are not affected by this command.

PG Command Example

This example shows how to indicate data paging is to be selected on the next invocation of the Indexed Access Method.

```
ENTER COMMAND (?): PG
DATA PAGING MARKED AS SELECTED
BECOMES EFFECTIVE ON NEXT LOAD OF $IAM
SEE INDEXED ACCESS METHOD GUIDE CONCERNING
REMOVAL OF PAGING MODULES FROM STORAGE.
```

```
ENTER COMMAND (?):
```



PP - DEFINE PAGING PARTITIONS

The PP command defines the amount of storage in each partition that the Indexed Access Method should reserve for paging. Storage is actually used for paging only when paging is active.

PP prompts you for the size of the paging area for each partition by displaying the partition number and current paging area size for that partition. Respond with a null entry (just press the Enter key) to retain that size. Enter a new size to change the space allocation. Sizes are displayed and entered in K bytes (1K = 1024), and should be entered as even numbers (multiple of 2K). If not, they are adjusted up to the next even number. The new sizes do not take effect until the next time the Indexed Access Method is loaded with paging active.

PP Command Example

This example retains the 10K paging area size in partition 3 and increases the paging area in partition 5 from 6K to 10K.

```
ENTER COMMAND (?): PP
PARTITION  CURRENT  NEW

   1          OK   :
   2          OK   :
   3          OK   : 40
   4          OK   :
   5          OK   :
   6          OK   :
   7          OK   :
   8          OK   :

PAGE AREA SIZE(S) RESET
BECOMES EFFECTIVE ON NEXT LOAD OF $IAM
TOTAL PAGE AREA SIZE IS  40K
SEE INDEXED ACCESS METHOD GUIDE CONCERNING
REMOVAL OF PAGING MODULES FROM STORAGE.
```

Notes:

1. The letter K is optional on input, and is assumed if missing.
2. The new total page area size is 40K and becomes effective on the next LOAD of \$IAM.

PS - GET PAGING STATISTICS

The PS command displays data paging information about the currently executing Indexed Access Method. It shows the total size of the paging area, and "hit" information for reads, writes and overall.

The Indexed Access Method increments a "hit" counter each time a referenced block is found in the paging area. It increments a "miss" counter each time a referenced block is not found in the paging area. The PS command displays these numbers, along with "hit percentages." Use the hit percentages to determine how efficiently the paging area is being used.

After the statistics are displayed, you have the option of resetting the counters to zero so that a new set of paging statistics can be gathered.

PS Command Example

Display the current paging statistic and reset them.

```
ENTER COMMAND (?): PS
FUNCTION      HITS      MISSES    HIT %
READ         45678     81205     36
WRITE        2450       0        100
OVERALL      48128     81205     37

RESET STATISTICS (Y/N)? Y
STATISTICS RESET
```

RE

RE - RESET PARAMETERS

RE resets the parameters set up by the SE command to their default values.

The following example shows a use of the RE command:

```
ENTER COMMAND (?): RE  
PARAMETERS RESET
```

RO - REORGANIZE INDEXED FILE

RO reorganizes a primary or secondary indexed file. It unloads an indexed file filled by insert activity into an empty indexed file and reorganizes the records to provide space for additional inserts.

This command requires two existing indexed files of the same type. Both the input file and the output file must be primary indexed files, or both must be secondary index files. Records are read sequentially from the input file using the Indexed Access Method GETSEQ request. The records are loaded into the output file in a manner similar to the initial load of the LO command.

All reserved and free space is retained as free space.

Reorganizing a secondary index does not reset the secondary key sequence numbers during the reorganization, because it does not use the primary file. The records are placed in another Indexed Access Method file without any modification within the individual records.

The output indexed file must have been defined by using the SE or DF commands before using the RO command. The SE Option 3 will format an output file like the original file, or \$VERIFY will show the number of records in the file so that you can set up an output file.

The record lengths of the two files need not be the same. Unloaded records are truncated or filled with binary zeroes if record lengths differ (see LO command). The key fields and key positions of the two files must be the same; however, the other file specifications (SE parameters) may differ.

INVOKING THE LOAD AND REORGANIZE FUNCTIONS FROM DF: You can invoke the LOAD or REORGANIZE functions directly from the DF command. If you invoke these functions, DF does not format the file because LOAD and REORGANIZE will do that, thus saving time. If you do not invoke the LOAD or REORGANIZE function, DF formats the file so you can load the file using an application program or the LO command.

Notes:

1. You can use the LOAD/REORGANIZE command later to load the file, if you do not invoke it from the DF command.
2. An application program cannot access an unformatted indexed file.
3. The prompt for the load/reorganize function occurs before the define step.

RO

The following example shows use of the RO command:

ENTER COMMAND (?): RO

REORG ACTIVE

ENTER INPUT DATASET (NAME,VOLUME): IAMFILE,EDX003

ENTER OUTPUT DATASET (NAME,VOLUME): IAMFIL2,EDX003

REORG IN PROCESS

END OF INPUT DATASET

100 RECORDS LOADED
REORG SUCCESSFUL

ENTER COMMAND (?): EN

SE - SET PARAMETERS

SE prompts you for parameters that determine the structure and size of the indexed file. An explanation of the SE command parameters follow and an example of each is included with the description.

The parameter values entered are saved by \$IAMUT1. This enables you to reuse the SE command to change one or more parameters without having to reenter all of them. The current values can be displayed by the DI command.

The SE command provides three methods of setting up an indexed file.

- Option 1** Significant Parameters - Enter a minimal set of SE parameters. The utility internally converts the smaller set to the complete set.
- Option 2** All Parameters - Enter the complete set of SE parameters.
- Option 3** Parameters from Existing Data Set - Use the set of SE parameters that were used previously to define an existing indexed file.

Note: Information which is common to all three options appears near the end of the SE description under "All Options" on page IG-182.

When you specify the SE command, you are prompted to select one of the options as shown in the following display.

```
SET FILE DEFINITION PARAMETERS
0 = EXIT
1 = SIGNIFICANT PARAMETERS
2 = ALL PARAMETERS
3 = PARAMETERS FROM EXISTING INDEXED DATASET
ENTER OPTION:
```

Option 1

Option 1 prompts for a minimal set of parameters. It issues a prompt to determine if a secondary index is being defined. If so, the secondary file name, key size, and key position are requested. If a primary file is being defined, different prompts are issued. \$IAMUT1 internally converts the option 1 parameters to option 2 parameters.

When the SE option 1 is invoked for the first time, the prompts and defaults values are as follows:

SE

SECONDARY INDEX (Y/N)? : N

	DEFAULT	NEW VALUE
RECORD SIZE	0:	80
KEY SIZE	0:	4
KEY POSITION	1:	
BLOCKING FACTOR (RECORDS PER BLOCK)	1:	
NUMBER OF BASE RECORDS	0:	20
ESTIMATED TOTAL RECORDS	24:	
TYPE OF INSERT ACTIVITY(C=CLUSTERED,R=RANDOM)	C:	

On subsequent invocations of the SE option 1, the defaults are taken from the parameter values since the last SE option 1 invocation. Option 1 and 3 values do not carry over to option 2.

The estimated total records value defaults to the last value, provided this value equals or exceeds the current base records. Otherwise it defaults to 1.2 times the current base records.

To set up a secondary index, enter the following:

```
SECONDARY INDEX (Y/N)? : Y
ENTER SECONDARY DATASET NAME (DS,VOL): FILE01,EDX002
SECONDARY KEY SIZE:10
SECONDARY KEY POSITION:36
```

Before you can define a secondary index, you must place an entry in the directory for the associated primary index file and the primary file must exist. The directory is searched to obtain the data set name and volume of the associated primary file which will then be used to compute the remainder of the secondary SE parameters.

Parameter Descriptions for Option 1

The attributes of the file are determined by the following SE command parameters:

RECORD SIZE: The length, in bytes, of each record in the file.

KEY SIZE: The length of the key to be used for this file. The minimum key length is 1. For primary files, the maximum key length is 254.


KEY POSITION: The position, in bytes, of the key within the record. The first byte of the record is position 1.

BLOCKING FACTOR (RECORDS PER BLOCK): The total number of records to be placed in an Indexed Access Method block. This value and the record size will be used to compute the actual Indexed Access Method block size, rounded up to the next 256-byte value. The rounding up action may increase the actual blocking factor.


NUMBER OF BASE RECORDS: The number of indexed record slots to be set up in the indexed file. These record slots can be loaded with data records by \$IAMUT1 or by a PUT request after either a LOAD or PROCESS request.

ESTIMATED TOTAL RECORDS: The total number of records you expect the indexed file to contain after insert processing activity.

TYPE OF INSERT ACTIVITY(C=CLUSTERED,R=RANDOM): Inserts are considered clustered if most of the inserts occur at only certain places in the file. The following diagram represents clustered inserts by vertically stacked bullets.



The next diagram represents random inserted records. Inserts are considered random if few or no points in the file have a concentration of activity; inserts are expected throughout the file.



SE

SECONDARY KEY SIZE: The length, in bytes, of the secondary key within the primary record. For secondary keys the maximum key length is 250.

SECONDARY KEY POSITION: The position, in bytes, of the secondary key within the primary record.

Option 2

The following list shows the default values for parameters when the SE command is invoked the first time (all values are decimal):

BASEREC	NULL
BLKSIZE	0
RECSIZE	0
KEYSIZE	0
KEYPOS	1
FREEREC	0
FREEBLK	0
RSVBLK	NULL
RSVIX	0
FPOOL	NULL
DELTHR	NULL
DYN	NULL

On subsequent invocations of the SE command, the option 2 defaults are taken from the parameter values set according to the last SE command, regardless of the option used. If the default value is acceptable, press the enter key when prompted for the parameter. If you wish to change the value for any parameter, enter the new value in response to the prompting message. The new value becomes the new default value for the current \$IAMUT1 session. The parameters for which a null can be specified are BASEREC, FREEREC, FREEBLK, RSVBLK, RSVIX, FPOOL, DELTHR, and DYN. To specify a null parameter after the original default has been modified, enter an ampersand (&) in response to the prompting message.

The following example shows a use of the SE command in establishing the size and structure of an indexed file.

PARAMETER	DEFAULT	NEW VALUE
BASEREC	NULL	:100
BLKSIZE	0	:256
RECSIZE	0	:80
KEYSIZE	0	:28
KEYPOS	1	:1
FREEREC	0	:1
FREEBLK	0	:10
RSVBLK	NULL	:
RSVIX	0	:
FPOOL	NULL	:
DELTHR	NULL	:
DYN	NULL	:

Following the response to the DYN parameter, the following list is displayed. The list shows the details of how the indexed file will be constructed using the parameters just entered.

TOTAL LOGICAL RECORDS/DATA BLOCK:	3
FULL RECORDS/DATA BLOCK:	2
INITIAL ALLOCATED DATA BLOCKS:	50
INDEX ENTRY SIZE:	32
TOTAL ENTRIES/INDEX BLOCK:	7
FREE ENTRIES/PIXB:	1
RESERVE ENTRIES/PIXB(BLOCKS):	0
FULL ENTRIES/PIXB:	6
RESERVE ENTRIES/SIXB:	0
FULL ENTRIES/SIXB:	7
DELETE THRESHOLD ENTRIES:	7
FREE POOL SIZE IN BLOCKS:	0
# OF INDEX BLOCKS AT LEVEL 1:	9
# OF INDEX BLOCKS AT LEVEL 2:	2
# OF INDEX BLOCKS AT LEVEL 3:	1
DATA SET SIZE IN EDX RECORDS:	73

SE

If a secondary file is being defined, the list of prompts is the same except for the following:

the reply to the prompt "SECONDARY INDEX (Y/N)?:" is Y

the secondary data set name is requested

the RECSIZE prompt is omitted; the Indexed Access Method computes the correct record size

```
SECONDARY INDEX (Y/N): Y
ENTER SECONDARY DATASET (NAME,VOLUME):
```

Parameter Descriptions for Option 2

The attributes of the file are determined by these SE command parameters:

BASEREC The estimated number of records to be initially loaded into the file in ascending key sequence. These records can be loaded by \$IAMUT1 or by a PUT request after either a LOAD or PROCESS request.

If DYN is not specified, BASEREC defaults to null, resulting in an error condition. In this case, specify BASEREC as a positive number.

If DYN is specified, BASEREC defaults to one.

BLKSIZE The length, in bytes, of blocks in the file. It must be a multiple of 256. The Indexed Access Method uses 16 bytes in each block for a header.

RECSIZE The length, in bytes, of each record in the file. Record length must not exceed block length minus 16.

KEYSIZE The length of the key to be used for this file. The minimum key length is 1. For primary files, the maximum key length is 254. For secondary a secondary index, the maximum key length is 250.

KEYPOS The position, in bytes, of the key within the record. The first byte of the record is position 1.

FREEREC The number of free records to be reserved in each block. It must be less than the number of records per block (block size minus 16, divided by record size). If not, an error message is

issued. The calculation is adjusted to ensure that there is at least one allocated record in the block; that is, there cannot be 100% free records.

FREEREC defaults to zero.

FREEBLK The percentage (0-99) of each cluster to reserve for free blocks. The percentage calculation result is rounded up so that at least one free block results. The calculation is adjusted to ensure that there is at least one allocated block in the cluster; that is, there cannot be 100% free blocks.

FREEBLK defaults to zero.

RSVBLK The percentage of the entries in each primary index block to reserve for cluster expansion. These reserved entries are used to point to new data blocks as they are taken from the free pool to expand the cluster. The result of the calculation is rounded up so that any non-zero specification indicates at least one reserved index entry. The calculation is adjusted to ensure that there is at least one allocated block in the cluster.

Enter a null character (&) for this prompt if you do not want initial reserved blocks and do not want the indexed access method to create reserved blocks as records are deleted and blocks become empty. Specify a value of zero for this prompt if you do not want initial reserved blocks but you do want the indexed access method to create reserved blocks as records are deleted and blocks become empty (See the DELTHR prompt).

Note that the sum of the FREEBLK and RSVBLK prompts must be less than 100 or an error message is issued. This value defaults to null if the DYN parameter is not specified. If the DYN parameter is specified, this value defaults to zero.

RSVIX The percentage (0-99) of the entries in each second level index block to reserve for use in case of cluster splits. A cluster split is required when there is no room to insert a new record in a cluster. Each cluster split uses one reserved entry of the second-level index block to create a new cluster with blocks from the free pool. The result of this calculation is rounded up so that any non-zero specification indicates at least one reserved index entry. The calculation is adjusted so that there is at least one unreserved entry in each second level index block. This value defaults to zero.

FPOOL The percentage (0-100) of the maximum possible free pool to allocate as determined by the RSVIX and RSVBLK parameters. The RSVBLK and RSVIX prompts result in a file structure set up to draw on the free pool for expansion.

If insertion activity is evenly distributed throughout the file, every reserve entry of every index block can be used. The number of blocks drawn from the free pool to support this unlikely condition is the maximum free pool size needed for the file. In more realistic cases, insertion activity is not evenly distributed throughout the file, so fewer free blocks are needed. The percentage specified here represents the evenness of the distribution of inserted records. Specify a large number (90, for example) if you expect insertions to be evenly distributed. Specify a small number (20, for example) if insertions are anticipated to be concentrated in specific key ranges.

If a null character (&) is specified for this prompt, a free pool is not created for this indexed file (you can use the DYN parameter to override this and create a free pool). If zero is specified, an empty free pool is created. Blocks can then be added to the free pool as records are deleted and blocks become empty (see the DELTHR prompt explanation). If you do not specify a null for this prompt, the RSVBLK must not be null and/or the RSVIX must be non-zero or an error is returned. Conversely, if the RSVBLK and/or RSVIX is non-zero, FPOOL must not be null or an error is returned.

The default for FPOOL is a null; no free pool is created.

DELTHR The percentage (0-99) of blocks to retain in the cluster as records are deleted and blocks made available. This is known as the delete threshold. When a block becomes empty, it is first determined if the block should be given up to the free pool by checking the response to this prompt. If the block is not given up to the free pool, it is retained in the cluster, either as a free block or as an active empty block. The result of this calculation is rounded up so that any non-zero specification indicates at least one block. The calculation is adjusted to ensure that the cluster always contains at least one block.

If the DELTHR parameter is specified as null (&) and DYN is not specified, DELTHR defaults to the number of allocated blocks in the cluster plus one half of the value calculated by the FREEBLK prompt. If the DELTHR parameter is specified as null and a value is specified for the DYN parameter, DELTHR defaults to zero.

DYN The number of blocks to be assigned to, or added to, the free pool. When DYN is used with other free pool parameters, the free pool size is calculated as specified by the FPOOL parameter plus the value specified for DYN.

If DYN is specified without the FPOOL parameter, the free pool is the number of blocks specified for DYN.

If DYN is specified, other parameters assume the following default values when specified as null:

```

BASEREC = 1
BLKSIZE = 0
RECSIZE = 0
KEYSIZE = 0
KEYPOS  = 1
FREEREC = 0
FREEBLK = 0
RSVBLK  = NULL
RSVIX   = 0
FPOOL   = NULL
DELTHR  = NULL

```

When you specify the number of blocks for the DYN parameter, remember that the Indexed Access Method can store several data records in a block, depending on the record size and block size you specify. Each block contains a 16 byte header. The number of records that can be contained in each block can be calculated by the following formula:

$$\text{Records per block} = (\text{BLKSIZE}-16)/\text{RECSIZE}$$

In the above calculation, use the integer quotient only; discard any remainder.

Blocks can be taken from the free pool for use as index blocks as well as for data blocks, so provide some extra blocks for these. A reasonable estimate of the number of index blocks required is 10%. Thus, if you know the number of data records you would like to add to the file, you can calculate the number of blocks to specify for the DYN parameter as follows:

$$\text{DYN} = \frac{\text{Number of records to insert} \times 1.1}{\text{Records per block}}$$

SE

Option 3

Option 3 issues a prompt to determine what existing file to obtain the parameters from. The parameters can be set exactly according to the parameters of the original file by replying Y to the appropriate prompt. Otherwise, the parameters will be set based on the current condition of the existing data set to reflect insert activity.

```
SECONDARY INDEX (Y/N)? : N
NAME OF EXISTING INDEXED DATA SET (NAME,VOLUME): IAMFILE,EDX003
NEW PARAMETERS EXACTLY SAME AS ORIGINAL PARAMETERS (Y/N) ? Y

DATA SETSIZE IN EDX RECORDS:          17
INDEXED ACCESS METHOD RETURN CODE:     -1
SYSTEM RETURN CODE:                   -1
```

All Options

For all three options, the prompts are followed by the option of entering the DF (define file) function directly from the SE command. This simplifies the file definition process. The prompt is as follows:

```
CREATE/DEFINE FILE (Y/N) ? Y
ENTER DATASET (NAME,VOLUME): FILE01,EDX003
NEW DATASET IS ALLOCATED
```

The immediate write-back option is then queried:

```
DO YOU WANT IMMEDIATE WRITE-BACK? Y
```

The next prompt allows you the option of invoking the load or reorganize functions as follows:

```
INVOKE LOAD(L), REORGANIZE(R) OR END(E) AFTER CURRENT FUNCTION? L
DEFINE IN PROGRESS
```

Size calculations are performed using the parameter values you specify.

SE

After the values are entered, the following is displayed showing the size and structure of the defined indexed file.

DATA SET SIZE IN EDX RECORDS:	17
INDEXED ACCESS METHOD RETURN CODE:	-1
SYSTEM RETURN CODE:	-1
CREATE/DEFINE FILE (Y/N)?:	

UN

UN - UNLOAD INDEXED FILE

UN unloads an indexed file to a sequential file. Records are read from the indexed file with the Indexed Access Method GETSEQ request and written into the sequential file with the EDL WRITE instruction. If a secondary indexed file is specified, the primary file will be unloaded in secondary key sequence.

You can unload a secondary index independent of its primary if you first use the UE subcommand of the DR command of \$IAMUT1 to set the independent indicator. You must turn the independent indicator off when the unload operation is completed.

The record lengths of the two files need not be the same. Unloaded records are truncated or padded with zeroes if the records lengths of the two data sets differ. For further detail, see the LO command.

Records are placed in the sequential file in ascending key sequence as indicated by the indexed file. Unloaded records can be blocked or unblocked. For a description of blocked and unblocked data sets, see "Blocked and Unblocked Sequential Data Sets" on page IG-164.

The UN command prompts you for the block size of the file to be unloaded. A null response or a value less than or equal to the record size causes the indexed file to be unloaded to an unblocked sequential file. The sequential file block size is calculated as the record size rounded up to the next 256-byte multiple value. If you want the file to be unloaded to a blocked sequential file, specify the actual block size value to the prompt "OUTPUT BLOCK SIZE". The record and block sizes of subsequent output sequential files are assumed to be the same as the initial output sequential file.

If the indexed file contains more records than are allocated in the sequential file, you are given the option to continue unloading to another sequential file. If you choose to continue unloading, you are prompted for the name of the file and volume to use to continue the unload operation. The unload operation continues, putting the records read from the indexed file into the new sequential file. If the end of the output file is reached and you choose not to continue, the unload operation ends.

Note: Do not specify the same file for input and output.

UN

The following example shows the use of the UN command to put 80-byte records into a blocked sequential file.

```
ENTER COMMAND (?): UN

UNLOAD ACTIVE
ENTER INPUT DATASET (NAME,VOLUME): EDXF02,AM4VOL
ENTER OUTPUT DATASET (NAME,VOLUME): SEQ01,EDX003
OUTPUT RECORD ASSUMED TO BE 80 BYTES. OK?: Y
ENTER OUTPUT BLOCK SIZE (NULL = UNBLOCKED): 256
UNLOAD IN PROCESS

END OF INPUT DATASET
  100 RECORDS UNLOADED
UNLOAD SUCCESSFUL
ENTER COMMAND (?): EN
```

\$IAMUT1 COMPLETION CODES

Completion Code	Condition
-1	Successful completion
7	Link module in use
8	Load error for \$IAM
12	Data set shut down
13	Module not included in load module
23	Get storage error - IACB
31	FCB WRITE error during IDEF processing, check system return code
32	Blocksize not multiple of 256
34	Data set is too small
36	Invalid block size during file definition processing
37	Invalid record size
38	Invalid index size
39	Record size greater than block size
40	Invalid number of free records
41	Invalid number of clusters
42	Invalid key size
43	Invalid reserve index value
44	Invalid reserve block value
45	Invalid free pool value
46	Invalid delete threshold value
47	Invalid free block value
48	Invalid number of base records
49	Invalid key position
50	Data set is already opened for exclusive use
51	Data set opened in load mode
52	Data set is opened, cannot be opened exclusively
54	Invalid block size during PROCESS or LOAD
55	Get storage for FCB error
56	FCB READ error
60	LOAD mode key is equal to or less than previous high key in data set
61	End of file
62	Duplicate key found

Note: For completion codes number 37 through 49, check your parameters for consistency.

Completion Code	Condition
65	Inconsistent free space parameters, RSVIX/RSVBLK or FPOOL/DYN were specified without one or the other
100	READ error, check system return code
101	WRITE error, check system return code
110	WRITE error - data set closed
201	Request failed because the primary file for this secondary could not be opened. Check system return code
210	Request failed because \$DISKUT3 could not be loaded
230	Directory read error from \$IAMDIR, verify that directory exists
231	\$IAMQCB not found, check sysgen for include of IIAMQCB
232	Directory open error for \$IAMDIR, verify that directory exists
233	Directory related primary request is a primary entry
234	Directory error - DSNAME,VOL not found in \$IAMDIR
235	Directory resource has not been requested
239	Directory write error. Refer to previously displayed message

CHAPTER 9. THE \$VERIFY UTILITY

\$VERIFY checks the validity of an indexed file and prints control block and free space information about the file on \$SYSPRTR.

This \$VERIFY description contains the following topics:

- \$VERIFY Functions
- Invoking \$VERIFY
- \$VERIFY Example
- \$VERIFY Messages
- \$VERIFY Storage Requirements

\$VERIFY FUNCTIONS

With \$VERIFY you can:

- Verify that all pointers in an indexed file are valid and that the records are in ascending sequence by key.
- Verify the contents of a secondary index against the primary file and report any discrepancies.
- Print a formatted File Control Block (FCB) listing, including the FCB Extension block. The FCB Extension block contains the original file definition parameters.

Note: The FCB Extension block does not exist and file definition parameters are not saved in the FCB for indexed files defined prior to version 1.2 of the Indexed Access Method. The reorganize (RO) \$IAMUT1 command can be used to reformat those files by adding an FCB Extension block to make use of all the \$VERIFY facilities.

- Print a report showing the distribution of free space in your file.
- Determine if any space is available for inserts.

INVOKING \$VERIFY

\$VERIFY can be invoked from either a terminal or a program coded in Event Driven Language. You supply the same input in either case. If you invoke \$VERIFY from a terminal, supply the input required in response to prompts. If you invoke \$VERIFY from a program, supply the input required as parameters passed to the program.

\$VERIFY INPUT

This section describes the input required to execute \$VERIFY.

name,volume Data set and volume names for the primary index file or secondary index to be processed. (Ensures that all chains within this data set are correct).

Option The type of processing you want \$VERIFY to do. The three options are:

Y - The FCB and the FCB Extension blocks are formatted and printed. The file is verified. A free space report is printed.

N - The FCB and the FCB Extension blocks are formatted and printed. The file is verified. No free space report is printed.

F - The FCB and the FCB Extension blocks are formatted and printed. No free space report is printed, but the '# OF AVAILABLE BLOCKS IN FREEPool' entry can be examined to determine if space is available for inserts; if the value is greater than zero (>0), space is available.

Cross verify option The type of check you want \$VERIFY to do between the primary index files and secondary indexes. The options are:

Y - a. If a primary index file was specified above as the data set name, this will check that all entries in the primary index file are in the secondary index.

b. If a secondary index was specified above as the data set name, this will check that all entries in the secondary index are in the associated primary indexed file.

N - Do not perform any cross verification.

secname,volume

Data set and volume names of the secondary index to be verified. Specify 'ALL' to verify all secondary indexes associated with the primary file.

INVOKING \$VERIFY FROM A TERMINAL

Load the \$VERIFY program as follows:

```
> $L $VERIFY
```

When \$VERIFY begins execution, you are prompted for the parameters described previously. A complete example of a \$VERIFY invocation from a terminal is shown under "\$VERIFY Example" on page IG-192.

INVOKING \$VERIFY FROM A PROGRAM

\$VERIFY can be invoked by EDL programs with the LOAD instruction. The only required parameter is the address of a 30-byte area that contains:

	Hex Displacement	Length (Bytes)
Data set name	0	8
Volume name	8	6
Detail listing request (Y, N, or F)	E	1
Secondary file cross verify (Y or N)	F	1
Secondary index file name	10	8
Secondary index file volume	18	6

The next example shows the use of \$VERIFY to verify a file named IAMFILE in the volume EDX002. A file verification and free space report are requested. The secondary file named SECIAM in the volume EDX002 is also verified.

```

EXAMPLE PROGRAM START
START EQU *
.
.
LOAD $VERIFY,PARMLIST,EVENT=VERIFY
WAIT VERIFY WAIT FOR POST COMPLETE
.
.
PROGSTOP
PARMLIST EQU *
DSNAME DC CL8'IAMFILE' INDEXED DATA SET NAME
VOLUME DC CL6'EDX002' VOLUME NAME
DETAIL DC CL1'Y' PROCESSING OPTION
SECONDRY DC CL1'Y' SECONDARY FILE VERIFICATION
SECDSN DC CL8'SECIAM' SECONDARY FILE NAME
SECVOL DC CL6'EDX002' SECONDARY FILE VOLUME
VERIFY ECB -1 EVENT CONTROL BLOCK
ENDPROG
END

```

\$VERIFY EXAMPLE

This section presents the input and output for an example run of \$VERIFY, along with descriptions of the material presented.

\$VERIFY is invoked from the terminal as follows:

```

[1] > $L $VERIFY
[2] INDEXED ACCESS METHOD FILE VERIFICATION PROGRAM ACTIVE
[3] ENTER NAME OF OUTPUT PRINTER. (BLANK = $SYSPRTR):
[4] (NAME,VOLUME): DPRIM1,EDXIAM
[5] DO YOU WANT DETAIL LISTING? (Y/N/F/?): Y
[6] DO YOU WISH TO VERIFY SECONDARY VS PRIMARY INDEXES (Y/N):N
[7] VERIFICATION COMPLETE, 0 ERROR(S) ENCOUNTERED
[8] $VERIFY ENDED

```

[1] In this example, the first line loads and executes \$VERIFY.

[2] The second line is printed by the program to indicate that execution has begun.

[3] This line allows you to direct the output to a particular printer or terminal. You can also press the Enter key without supplying a device name and the output will be printed on \$SYSPRTR.

[4] In the fourth line, the program prompts for the data set name and volume of the indexed file to be referenced by the program. In this example the reply indicates that the data set is DPRIM1, located on volume EDXIAM.

[5] In the fifth line, the program prompts for the amount of detail to be provided as output. The response of Y indicates that maximum detail is to be provided.

[6] In the sixth line, the program prompts for verification of secondary indexes. The response of N indicates that secondary indexes are not to be verified. As the program executes, it provides output to the printer, as shown in the example outputs that follow.

[7] Finally, messages are displayed to indicate the number of errors found. [8] This information message is provided stating that the program has ended.

FCB REPORT

The first page of the example output from \$VERIFY follows. This page is always printed.

VERIFY REPORT. FILE = DPRIM1 , VOLUME = EDXIAM

FLAG1 : FILE FILE
LOADED TYPE
Y 1 (0=PRPQ, 1=PP)

KEY SIZE = 6
KEY POSITION = 1
BLOCK SIZE = 256
RECORD SIZE = 60
INDEX ENTRY SIZE = 10
RBN OF HIGH LEVEL INDEX BLOCK IN USE = 2
RBN OF LAST DATA BLOCK IN USE = 786
RBN OF FIRST DATA BLOCK IN USE = 6
TOTAL RECORDS PER DATA BLOCK = 4
TOTAL ENTRIES PER INDEX BLOCK = 24
LOAD POINT VALUE FOR A DATA BLOCK = 4
LOAD POINT VALUE FOR AN INDEX BLOCK = 24

FLAG2 : IMMEDIATE SECONDARY FILE
WRITE-BACK INDEX FILE FORMATTED
N EN Y

VERSION NUMBER = 2.0
DELETE THRESHHOLD (RECORDS) = 0
OF AVAILABLE BLOCKS IN FREEPool 30
RBN OF 1ST FREE POOL BLOCK = 787
RBN OF HIGHEST LOGICAL INDEX BLOCK = 2
LEVEL OF HIGHEST INDEX BLOCK IN USE= 3
CURRENT NO. OF RECORDS IN FILE = 3000

The preceding page of the example report is interpreted as follows:

The first line shows the data set name and volume.

FLAG1: These three lines show the significant bits of the first flag byte in the FCB. The first two of the three lines are a heading. The third line shows the bit value (1 = on and 0 = off or Y = on and N = off). The headings are defined as follows:

FILE LOADED: Data set has been loaded flag. This flag is set when any record has been successfully loaded into the file in load mode.

FILE TYPE: This flag indicates whether the indexed file was created with the Realtime Programming System Indexed Access Method PRPQ (bit=0) or either the Event Driven Executive or Realtime Programming System Indexed Access Method Program Product (bit=1).

KEY SIZE: Shows the size of the key in bytes.

KEY POSITION: Shows the byte displacement of the key from the start of the record.

BLOCK SIZE: Shows the byte length of blocks in the file.

RECORD SIZE: Shows the byte length of records in the file.

INDEX ENTRY SIZE: Shows the number of bytes in each index entry. This length should be the key length plus 4, rounded up to a multiple of two bytes.

RBN OF HIGH LEVEL INDEX BLOCK IN USE: Shows which index block is to be used as the starting point when the index is to be searched.

RBN OF LAST DATA BLOCK IN USE: Points to the last logical data block in the file which has been used.

RBN OF FIRST DATA BLOCK IN USE: Points to the first logical data block in the file which has been used. It is used as the starting point when a sequential read operation is begun with no key specified.

TOTAL RECORDS PER DATA BLOCK: Shows how many data records can be contained in a data block.

TOTAL ENTRIES PER INDEX BLOCK: Shows how many index entries can be contained in an index block.

LOAD POINT VALUE FOR A DATA BLOCK: The number of records that can be placed in each data block while in load mode. This value is calculated at file definition time to provide the requested number of free records.

LOAD POINT VALUE FOR AN INDEX BLOCK: The number of data blocks in each cluster to be used while in load mode. This value is calculated at file definition time to provide the space requested by the RSVBLK, RSVIX and FREEBLK parameters.

FLAG2: Another byte of flags described by a pair of lines: a heading line followed by a data line. The heading has the following meaning:

IMMEDIATE WRITE-BACK: Immediate write back flag. If set (Y), this flag indicates that the immediate write back option was specified when the indexed file was defined.

SECONDARY INDEX FILE: A Y indicates that this is a secondary file. N indicates that this is a primary file.

FILE FORMATTED: Y indicates that the file has been formatted. N indicates that only the parameters have been specified and the file allocated. The file has not been formatted.

VERSION NUMBER: Shows the version number and modification level of the Indexed Access Method that was used to define the indexed file.

DELETE THRESHHOLD (RECORDS): Indicates the number of data blocks to retain in each cluster as records are deleted and blocks become empty. This value is calculated when the file is defined and is based on the DELTHR parameter.

OF AVAILABLE BLOCKS IN FREEPOOL: The number of available blocks in the free pool. This count is updated as blocks are taken from or returned to the free pool.

RBN OF 1ST FREE POOL BLOCK: Points to the last block which was put in the free pool (which is the next block to be taken from the free pool).

RBN OF HIGHEST LOGICAL INDEX BLOCK: Points to the logical top of the index. In some cases (if the file has not been completely loaded), this RBN might not agree with the RBN OF HIGHEST LEVEL INDEX BLOCK IN USE. If it does not agree, then the file is structured with index blocks that are not yet needed because the file does not contain enough records.

LEVEL OF HIGHEST INDEX BLOCK IN USE: Indicates how many levels of the index are currently in use.

CURRENT NO. OF RECORDS IN FILE: The current number of records that are now contained in the file.

FCB EXTENSION REPORT

The second page of the example output from \$VERIFY follows. This page is always printed.

This information is obtained from the FCB Extension block and shows the parameters that were specified when the file was defined. Some information (BLKSIZE, RECSIZE, KEYSIZE, KEYPOS) is duplicated on the FCB and FCB Extension report because it is contained in both control blocks. The values should correspond with each other. The word NULL for the value of a parameter indicates that no value was specified when the file was defined.

VERIFY REPORT. FILE = DPRIM1 , VOLUME = EDXIAM

INDEX FILE DEFINED WITH THESE PARAMETERS:

BASEREC=	3000
BLKSIZE=	256
RECSIZE=	60
KEYSIZE=	6
KEYPOS=	1
FREEREC=	0
FREEBLK=	0
RSVBLK=	NULL
RSVIX=	0
FPOOL=	NULL
DELTHR=	NULL
DYN=	30

Note: The parameters are the file definition parameters that were specified using the SE command of the \$IAMUT1 utility when the file was defined.

FREE SPACE REPORT

A free space report of the example output from \$VERIFY follows. The free space report is printed only if the \$VERIFY option is specified as Y.

VERIFY REPORT. FILE = XMPL1 , VOLUME = EDX002

RBN	LVL	TOTAL ENTRIES	USED ENTRIES	UNUSED ENTRIES	RESERVE ENTRIES	FREE BLOCKS	AVAILABLE RECORD SLOTS	HIGH KEY (FIRST 20 CHAR.)
-----	-----	---------------	--------------	----------------	-----------------	-------------	------------------------	---------------------------

2	3	24	2	0	22	0	--	143949
3	2	24	24	0	16	0	--	130536
4	2	24	8	0	16	0	--	143949
5	1	24	24	0	0	0	0	044932
30	1	24	24	0	0	0	0	046750
55	1	24	24	0	0	0	0	048655
80	1	24	24	0	0	0	0	050527
105	1	24	24	0	0	0	0	052392
130	1	24	24	0	0	0	0	054225
155	1	24	24	0	0	0	0	056075
180	1	24	24	0	0	0	0	057930
205	1	24	24	0	0	0	0	059829
230	1	24	24	0	0	0	0	061640
255	1	24	24	0	0	0	0	063548
280	1	24	24	0	0	0	0	065389
305	1	24	24	0	0	0	0	067297
330	1	24	24	0	0	0	0	069166
355	1	24	24	0	0	0	0	071029
380	1	24	24	0	0	0	0	072887
405	1	24	24	0	0	0	0	074731
430	1	24	24	0	0	0	0	076586
455	1	24	24	0	0	0	0	078441
480	1	24	24	0	0	0	0	080329
505	1	24	24	0	0	0	0	082175
530	1	24	24	0	0	0	0	084006
555	1	24	24	0	0	0	0	085861
580	1	24	24	0	0	0	0	130536
605	1	24	24	0	0	0	0	132395
630	1	24	24	0	0	0	0	134205
655	1	24	24	0	0	0	0	136097
680	1	24	24	0	0	0	0	137929
705	1	24	24	0	0	0	0	139815
730	1	24	24	0	0	0	0	141655
755	1	24	24	0	0	0	0	143523
780	1	24	6	0	18	0	0	143949

VERIFICATION COMPLETE, 0 ERROR(S) ENCOUNTERED

In this report, each printed line represents an index block. The columns have the following meanings:

RBN: The relative block number within the indexed file, based on the block size specified when the file was defined. The first block in the file is relative block number zero.

LVL: The level of the index block analyzed. Lowest level (PIXB) is 1, second level (SIXB) is 2, etc.

TOTAL ENTRIES: The maximum number of index entries that can fit in an index block.

USED ENTRIES: The number of entries used in this index block.

UNUSED ENTRIES: The number of entries in the index block which are neither used nor reserved.

RESERVE ENTRIES: The number of reserve entries in this index block. This number represents the number of new index blocks that can be obtained from the free pool for creation of new blocks, provided there are enough blocks remaining in the free pool.

FREE BLOCKS: The number of free blocks associated with this index block.

AVAILABLE RECORD SLOTS: The maximum number of records that can be inserted into this cluster without obtaining blocks from the free pool.

HIGHEST KEY IN BLOCK: The first 20 bytes of the highest key in the block.

\$VERIFY MESSAGES

As \$VERIFY executes, any errors encountered result in an error message being written describing the type of error and where the error occurred.

FILE ERROR MESSAGES

The following messages indicate that the indexed file contains errors:

BLOCKS IN FREEPOOL CHAIN DOES NOT MATCH FREE POOL COUNT IN FCB.

BLOCK OUT OF SEQUENCE. RBN ____.

HIGH KEY IN RBN ____ DOES NOT MATCH INDEX ENTRY IN RBN ____.

POINTERS IN HEADER OF HIGH INDEX BLOCK ARE NOT ZERO.

RBN ____ CONTAINS INVALID UPWARD POINTER.

RBN ____ CONTAINS INVALID BACKWARD POINTER.

RBN ____ CONTAINS INVALID FORWARD POINTER.

RBN ____ IS IN FREEPOOL CHAIN, BUT IS NOT A VALID FREEPOOL BLOCK.

RECORD OUT OF SEQUENCE NEAR RBN ____.

RECORD MATCH NOT FOUND FOR SEC INDEX.

PRIMARY=

SECONDARY=

If any of these messages are printed, the indexed file has at least one error.

Possible sources of the error include:

- The data set is not an indexed file
- Data in the file has been inadvertently destroyed
- Secondary index is not auto-update
- The Indexed Access Method has a program error

ERROR RECOVERY PROCEDURE

If any of the \$VERIFY file error messages are printed, use the following procedure:

- Dump the file or portion of the file which \$VERIFY indicated has errors.
- Attempt to reorganize the file with the \$IAMUT1 utility RO command.

- If reorganization fails, submit an APAR, including the file dump.
- Secondary indexes may need to be regenerated. Invoke \$VERIFY for each of the secondaries to determine if they are error free. If errors are indicated rebuild the index from the primary data sets after the problem has been corrected.

OTHER MESSAGES

Any other messages that are generated by \$VERIFY can be found in the Messages and Codes.

\$VERIFY STORAGE REQUIREMENTS

Working storage space is needed required for \$VERIFY and the amount required varies, depending on the maximum number of blocks at the SIXB level and the block size of the file.

USING DEFAULT WORKING STORAGE REQUIREMENTS

The default working storage specification is 4K bytes. For a file with a block size of 256, this default is sufficient to handle up to 896 blocks at the SIXB level. The larger the block size of the file, the fewer the maximum number of SIXBs that can be processed.

The following formula can be used to calculate the maximum number of blocks at the SIXB level that \$VERIFY can process, given the block size of the indexed file:

$$NS = (4096 - (2 * BLKSIZE)) / 4$$

NS is the number of blocks at the SIXB level
BLKSIZE is the block size of the indexed file

MODIFYING WORKING STORAGE REQUIREMENTS

The default working storage allocation is intended to satisfy the requirements of most indexed files. It may be necessary or desirable to modify the amount of working storage space available to \$VERIFY.

The following formula can be used to calculate the amount of working storage required to process a file with a given block size and number of blocks at the SIXB level.

$$DS = (4 * NS) + (2 * BLKSIZE)$$

Where:

DS is the amount of dynamic storage required

NS is the number of blocks at the SIXB level

BLKSIZE is block size of the indexed file

The number of SIXBs in a file can be determined by examining the free space report.

You can override the default working storage size at load time (if loaded by a program), or with the SS command of the \$DISKUT2 utility.

SUMMARY

\$VERIFY requires a variable amount of working storage which defaults to 4K bytes. Increase the working storage size if \$VERIFY runs out of space during execution.

Decrease the working storage size if the number of SIXBs is significantly less than that supported by the default working storage allocation (896 with a block size of 256) and your available storage is limited.

CHAPTER 10. STORAGE AND PERFORMANCE CONSIDERATIONS

This chapter describes the storage required for the Indexed Access Method and suggestions for improving performance. The main topics are:

- Determining Storage Requirements
- Data Paging
- Other Performance Considerations

DETERMINING STORAGE REQUIREMENTS

The minimum amount of storage required by the Indexed Access Method is dependent upon the package you choose to install, plus the link module and any error exit routine you may have written. The approximate sizes of the available packages are included here for planning purposes.

THE INDEXED ACCESS METHOD PACKAGES

The Indexed Access Method program product is shipped with four packages:

- \$IAM
- \$IAMRS
- \$IAMNP
- \$IAMRSNP

You select the particular package to install on your system which meets your requirements for function, storage, and performance. The individual packages are described below:

1. \$IAM - (18K), a full function Indexed Access Method package using overlay structure. It is expected to satisfy the needs of most users.
2. \$IAMRS - (27K), a full function Indexed Access Method as a fully resident package. It requires more storage than \$IAM, but offers maximum performance.
3. \$IAMNP - (15K), this package is similar to \$IAM (using an overlay structure) but does not include data paging. It is designed for users who have severe storage limitations.
4. \$IAMRSNP - (24K), this package is similar to \$IAMRS (fully resident) but does not include data paging. This package provides the perform-

ance of a resident system but is intended for users who do not have sufficient storage to take advantage of the data paging feature.

Notes:

1. The storage values above do not include Indexed Access Method control blocks, the central buffer (minimum of 2 X block size), and secondary index update buffers (minimum of 2 X record size).
2. To find the exact size of your Indexed Access Method package, load \$IAM with the operator command \$L. A message will be displayed about the loaded program. The number, followed by the letter P, indicates the size of the program in 256-byte pages. Multiplying this number by 256 yields the size in bytes of \$IAM, including control blocks, work areas and buffers.

INDEXED ACCESS METHOD STORAGE ENVIRONMENT

A single copy of the Indexed Access Method load module \$IAM serves the entire system.

Figure 12 on page IG-205 shows the components of the Indexed Access Method, and their relationship to the operating system.

The Indexed Access Method control blocks, buffers and programs are contained in a single module, which can be loaded in any partition (but only one copy on the system).

Application programs in any partitions (including the partition containing the Indexed Access Method) can invoke Indexed Access Method services using the IBM supplied link module, which must be included in the application program.

If the data paging feature of the Indexed Access Method is active, it uses storage in the partition(s) you select for performance improvement. This storage is in the form of a load module, \$IAMSTGM.

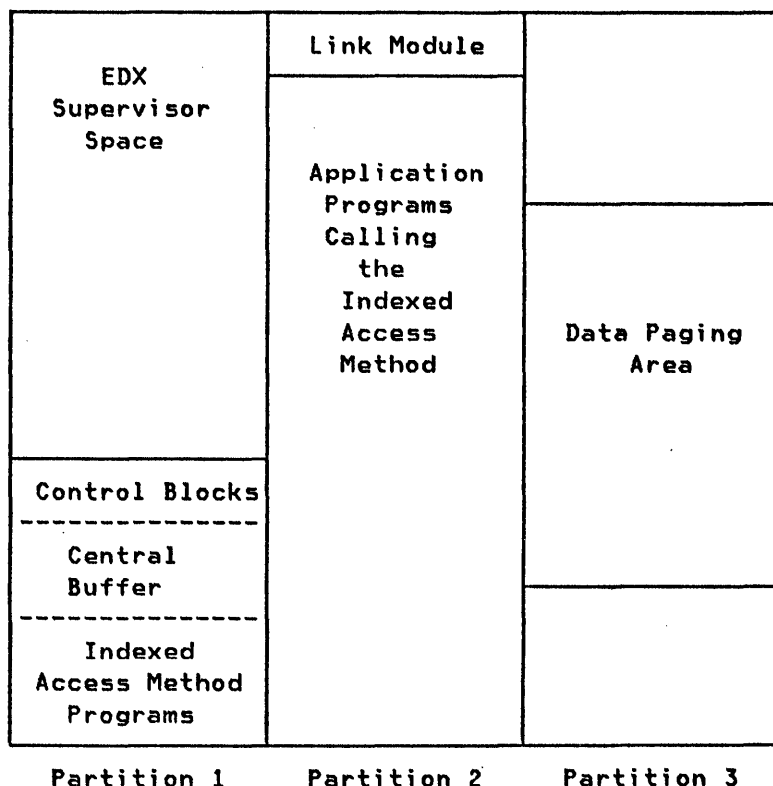


Figure 12. Indexed Access Method Storage Environment

Because \$IAM is loaded automatically when the first Indexed Access Method request is issued, it does not need to be explicitly loaded before being used by any program. When loaded automatically on the first Indexed Access Method request, \$IAM is loaded into partition 1 if enough storage is available there. If not, attempts are made to load \$IAM into successively higher numbered partitions until space is found or no more partitions are available. Once loaded, the Indexed Access Method remains in storage until cancelled with the \$C operator command.

The Indexed Access Method can also be loaded manually by using the \$L operator command or automatically at IPL time through the automatic initialization capability (refer to the System Guide for a description of the automatic initialization capability). \$IAM can be loaded into any partition, including partition one. It can be invoked (through the link module) from application program in any partition.

PERFORMANCE

Performance can be improved by various factors and the performance will be different for each application. One performance consideration has been described previously, the resident Indexed Access Method packages \$IAMRS and \$IAMRSNP. Another supplied performance feature is data paging.

DATA PAGING

Data paging is a performance feature that uses main storage space for a paging area (a cache) to improve the performance of the Indexed Access Method. This paging area retains recently used index and data blocks which have been retrieved for processing. As blocks are read from an indexed file, they are retained in the paging area on the assumption that they will probably be requested again. When a block is requested again, if it is in the paging area, no I/O operation is required; the block is moved directly into the central buffer.

The paging area is divided into 2K-byte (2048-byte) pages. Each indexed file can also be thought of as being divided into 2K-byte pages. When data is read from the file, a 2K-byte page is read and saved in the paging area. When data is written to the file, only the modified block (not the 2K-byte page) is written.

When the paging area becomes full, pages are overlaid according to a least-recently-used algorithm. The Indexed Access Method data paging algorithm handles direct access records differently from the way it handles sequential access records.

SEQUENTIAL ACCESS AND DATA PAGING: All of the pages in the page area can be used for direct access. However, because sequential access can cause the page area to be flushed out (negating the advantages of data paging), only 25% of the pages are set aside for use in sequential mode. Therefore, pages referenced in sequential mode will only use a small portion of the page area. This causes the pages to tend to preempt themselves instead of flushing out the page area.

REMOVAL OF STORAGE MODULES: The data paging area is obtained by loading a copy of \$IAMSTGM into one or more partitions. Each copy of \$IAMSTGM remains in storage, even if you cancel \$IAM. Cancelling \$IAM is not recommended unless you have ascertained that no files are currently open and no requests are about to be issued. If you have cancelled \$IAM you can use the \$C \$IAMSTGM operator command to remove the data paging storage module from each partition. \$IAMSTGM should never be cancelled until you have first cancelled \$IAM.

Adjusting the Size of the Paging Area

Because every application is different you should not regard any information relative to the following described example as being directly applicable to your application. However, the general principles should apply to most applications.

Figure 13 on page IG-208 shows the effect of various data paging area sizes on the percentage of times a requested block was in the paging area ("Hit Ratio") and the resultant performance (response time indicator) for one application. The data was acquired by measuring the performance, and printing data paging statistic, while the application was running. The total size of all indexed files being accessed during the run was 36592

sectors (9.3M-bytes). It must be stressed that this is only one application, and your application may not behave in the same manner.

The three variables considered in data paging described in this example are:

- storage size dedicated to data paging
- the percentage of times that the block requested is in the paging area ("Hit Ratio")
- read/write ratio

STORAGE SIZE: The figure shows general trends for various storage sizes. Note that there is a minimum amount of storage which can provide a benefit. In this example the minimum storage to acquire a performance improvement is approximately 20k-bytes. This is because the data paging algorithms in the Indexed Access Method require a certain amount of processing, which is additional overhead. Your application may have a different minimum. If you cannot supply enough storage to provide a benefit, you are better off not to use data paging. Within certain limits, the more storage you supply, the better the performance. However, there are optimal minimum and maximum limitations. Figure 13 on page IG-208 shows that, for this example application, the minimum amount is about 20k-bytes.

The optimal maximum amount of storage, beyond which the benefit of using more storage becomes less pronounced, is about 70k-bytes for the example shown in Figure 13 on page IG-208. You must determine, based on your own storage/performance tradeoff requirements, how much storage to dedicate to data paging for the performance improvement you receive. Larger files require a proportionately larger paging area to attain the same hit ratio.

THE "HIT RATIO": The values shown at the left side of Figure 13 on page IG-208 is called a "Hit Ratio". This ratio is a percentage of how often an index block or data block requested is already in the paging area. Most applications tend to concentrate activity in a few areas of the file for a time, then move on to other areas of the file. These applications can use data paging to good advantage because there is a probability that the data being requested has been recently requested.

If your application references data in a completely random manner, data paging will be less efficient. Random applications result in a smaller hit ratio for a given paging area size than applications that concentrate on certain areas of the file. Therefore, larger paging area is required to obtain the same hit percentage.

THE READ/WRITE RATIO: The data paging function is optimized for read operations. In order to insure file integrity, write operations cause a write-through to the file. This means that there is no benefit in using data paging for write operations. In fact, due to paging overhead, write operations are less efficient with data paging than without data paging.

The higher your ratio of reads to writes, the more efficiently the data paging algorithm works, thus the better your performance improvement. In

the example shown in Figure 13 on page IG-208, 80% of the requests were reads, 20% were writes.

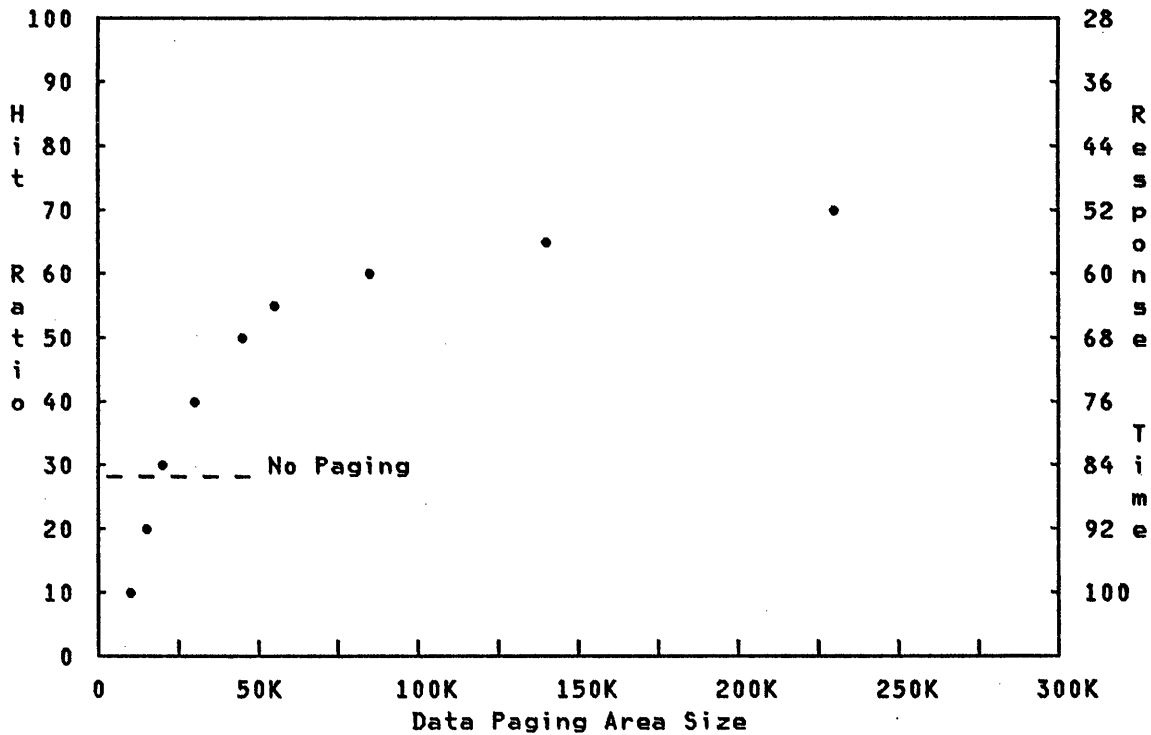


Figure 13. Plot of Data Paging Area Sizes. This graph shows how the size of the data paging area (shown across the bottom) affects the hit ratio (shown on the left margin), are the results in the response time (shown on the right margin). The unit of time for the response time scale is not given because it is application dependent. For this application, a hit ratio of at least 28% (which can be achieved with a paging area size of about 20K) is required to attain performance equal to that without data paging active. This is due to data paging processing overhead. Also note that a paging area size of greater than about 70K provides relatively little response improvement for the amount of storage dedicated.

Using Data Paging

The Indexed Access Method is distributed with the paging area size set to zero; therefore, the data paging function is not enabled. To use paging, use the \$IAMUT1 PP command to set the paging area size for each partition and the \$IAMUT1 PG command to activate paging.

When \$IAM is loaded, the loader attempts to obtain storage in the requested partition. When storage is requested in a particular partition to activate paging, you are informed of the results with appropriate messages. The messages returned to inform you of the paging status are written to the \$SYSLOG device. If \$SYSLOG is not available, the messages are

written to \$SYSPRTR device. Following are the conditions which can result:

1. Data paging is successfully initialized. The storage you requested or the default amount of storage required for data paging is available. The following message is displayed:
"DATA PAGING ACTIVE."
2. Data paging is unsuccessful.
 - a. If you have attempted to activate paging and you requested zero for the storage amount, or the minimum amount of storage necessary for paging is not available in the partition you specified, the following message is displayed:
"NOT ENOUGH STORAGE AVAILABLE FOR DATA PAGING."
"DATA PAGING NOT ACTIVE."
 - b. If you have requested more storage for paging than is available in the partition you specified, the following message is displayed
LOAD FAILED FOR \$IAMSTGM RC=xxx, PTN= Y, SIZE= zz

where: xxx represents the return code from the LOAD instruction
 y represents the partition number requested
 zz represents the size in 2048-bytes of storage
 you requested

OTHER PERFORMANCE CONSIDERATIONS

Following is a list of subjects followed by some ideas you might use to affect the performance of your application:

- Looking at the File Structure
- Controlling the File size
- Reducing the Number of Index Levels
- Increasing the Buffer Size
- Avoiding Resource Contention

LOOKING AT THE FILE STRUCTURE: Performance of the Indexed Access Method is primarily determined by the structure of the indexed file being used. This structure is determined by parameters you specify when you create the file. The best performance from an indexed file is attained when the file structure is well planned and the free pool is rarely used, if it exists at all. For descriptions of the file parameters, see "Chapter 12. Installing the Indexed Access Method." For examples of the affects of parameter values, see "Chapter 3. Defining Primary Index Files."

Use the \$IAMUT1 utility to see the effects of the various parameters on the file structure.

FILE SIZE: A large file spans more cylinders of the direct access device, so the average seek to get the record you want is longer. Splitting files into smaller files according to application type, or moving seldom used records to a "history file" might be viable solutions for file size reductions.

If your records contain unused or unnecessary fields, delete those fields and reduce your record length before defining and loading your file. The Sort/Merge Program Product contains facilities to accomplish this while sorting your records by key.

REDUCING THE NUMBER OF INDEX LEVELS: A file with many index levels requires more accesses to get to the desired data record, thus degrading performance. Factors which influence the number of index levels are:

- Number of records in file - see "File Size" previously described
- Amount and type of free space - see "File Structure" previously described
- Block size - when defining your indexed file, remember larger block sizes usually require fewer I/O operations
- Key size - shorter keys are more efficient than long keys. If only a portion of your key field provides uniqueness, set your key position and key length to that portion of the field when you define the file.

INCREASING THE BUFFER SIZE: The buffers required for I/O operations for all Indexed Access Method requests throughout the system are taken from a single buffer pool. The size can be changed at any time (to become effective during the next load of \$IAM) as described under "RE - Reset Parameters" on page IG-170. If you provide a large buffer when you install the Indexed Access Method, it is more likely that blocks (especially high-level index blocks) needed are already in storage and need not be recalled from the file.

AVOIDING RESOURCE CONTENTION: Application programs that use the Indexed Access Method are executed the same as other application programs. Because the Indexed Access Method and the indexed files are resources available to all tasks, delays can occur. When more than one task uses the Indexed Access Method, contention can occur between tasks for any of the following resources:

- An entire indexed file
- An index block in the file
- A data block in the file
- A data record in the file
- Buffer space from the system buffer pool

For example, during the execution of a request from task A, some buffer space is required and an index block, data block, or record is locked (made unavailable to other requests). A request from task B requires more buffer space than is available or attempts to retrieve a block or record

that was locked by task A. Task B must wait until the required resource becomes available.

Resources required by the Indexed Access Method are allocated only for the duration of a request except under the following circumstances:

- During an update, when control returns to the task after a GET or GETSEQ for update, the subject record is locked. The lock is released when the update is completed with a PUTUP, PUTDE, RELEASE, or DISCONN.
- During sequential processing, when control returns to the task after a GETSEQ, the block containing the subject record is locked and held in the buffer.

Subsequent GETSEQ requests pick up records directly from the buffer. When a GET requires a record from the next block, the current block and buffer are released. Pending requests for a buffer area are satisfied and the next block is locked and held in the buffer. Except for momentary release of the buffer area between blocks, a block is locked while it is being processed. Processing is terminated by an end-of-data condition, an ENDSEQ request, a DISCONN request, or an error condition.

Use the following guidelines to avoid resource contention:

- Disconnect all indexed files before task termination. The DISCONN request releases locked records or blocks and writes records that have not already been written.
- Use conditional requests whenever possible so that your application can be productive while a resource is unavailable.
- Try to schedule applications so that they do not execute at the same time.
- If a file is used for "read only" by more than one application, consider multiple copies of the file using unique file names.
- With multiple Indexed Access Method applications, use direct access to retrieve a group of records. A suggested method is the following:
 1. Retrieve the first record by key.
 2. Extract the key from the record and save it for the next retrieval.
 3. Retrieve the next record using the saved key and a greater than key relational operator (GT or UPGT).
 4. Repeat the second and third steps until processing is complete.

Secondary Index Functions

Using secondary indexes affects the performance of the Indexed Access Method. Some of those reasons are described here.

DIRECT RETRIEVAL: Direct retrievals are somewhat slower when using a secondary index because of the extra accesses required to retrieve the data record from the primary file.

SEQUENTIAL RETRIEVAL: Sequential Retrievals are slower when using a secondary index because the records are returned in order by secondary key. The primary file containing the data records is in order by the primary key. Therefore, the records are not stored in the same sequence that they are retrieved. This requires random accesses to obtain the records.

RECORD INSERTS: Record inserts are slower if any associated secondary indexes have the auto-update indicator on. A new record must be inserted into each auto-update secondary index, as well as the primary, whether the original insert was a primary or a secondary.

RECORD DELETE: Record deletes are slower for the same reason as for inserts; records must be deleted from secondary indexes that have the auto-update indicator on. However, the impact for deletion is more severe than for insertion. This is because a search is required when multiple records have the same value for their secondary key as the record being deleted. The group of records having the same key must be sequentially searched until the record with the required primary key is found. This time could be quite significant if you have large groups of duplicate keys.

RECORD UPDATE: Record updates that modify the secondary key must also update any associated secondary index which has the auto-update indicator on. The secondary index is updated by deleting the old key and inserting the new key.

DATA RECORD MOVEMENT: Each record in a secondary index contains a pointer to the RBN where the record is located in the primary index file. If a data record has been moved, due to insert/delete activity in nearby areas of the primary file, the RBN in the secondary index record will be wrong. When the affected data record is retrieved through the secondary index, the error is detected. A full retrieval is then performed, using the primary key to obtain the data record. The RBN in the secondary index record is then updated for the benefit of future retrievals. This activity will affect the performance.

\$VERIFY PERFORMANCE: The \$VERIFY performance will be slower when the primary file being verified has a secondary index with large numbers of duplicate secondary keys. This is because the entire group of duplicate keys must be searched for the proper record. Because \$VERIFY retrieves all records in the file, these impacts accumulate and the total execution time can be longer than expected.

CHAPTER 11. ERROR RECOVERY

This chapter describes how to handle Indexed Access Method errors and how to diagnose application program errors.

HANDLING ERRORS

All Indexed Access Method requests return a code in the task code word of the Task Control Block (TCB). The task code word is the same name as the task name. The return code reflects the condition of the requested function. Return codes are grouped in the following categories:

- -1 - Successful completion
- Positive - Error
- Negative (other than -1) - Warning

ERROR EXIT FACILITIES

There are three types of error exits for your application:

- Task error exit, provided by the supervisor
- Error exit, provided by the Indexed Access Method
- The task error exit used by the Indexed Access Method itself in case of an error

Task Error Exit

You can specify a task error exit routine that will receive control if your application program causes a soft exception or if a machine check occurs during the execution of your application.

Because your application may have requests pending (for example, a record is being held for update or a file is being processed sequentially), you should issue a DISCONN request before terminating your application. The task error exit allows you to release records, disconnect from any file you are connected to, and make your resources available to other applications. Use of the task error exit facility helps to ensure data integrity and allows proper termination or continuation of your application.

Implementing the task error exit facility is described in the System Guide.

Error Exit

In PROCESS and LOAD requests, the address of an error exit routine can be specified by the ERREXIT parameter. If specified, this routine is executed whenever an Indexed Access Method request terminates with a positive return code.

If the exit routine is not specified, the next sequential instruction after the request is executed regardless of the value of the return code.

\$IAM Task Error Exit

The Indexed Access Method itself has a task error exit. If this error exit is given control by the supervisor, it writes these messages to the \$SYSLOG device:

```
$IAM HAS INCURRED A SEVERE ERROR
$IAMCENTRAL BUFFER ADDRESS IS xxxx - PARTITION n
PSW  LSB
yyyy zzzz zzzz zzzz zzzz zzzz zzzz zzzz zzzz zzzz zzzz zzzz
```

Where xxxx is the address, and n is the partition containing \$IAM.

The PSW (yyyy), and LSB contents (zzzz), are also listed. For an explanation of the PSW (program status word) and the LSB (level status block), refer to the Operator's Reference.

\$IAM then goes into an unrecoverable wait and will not process any access requests. You can dump the central buffer with the \$D system command and take appropriate action to quiesce your application (refer to the Operator's Reference for a description of the \$D command).

You can use the recovery and backup procedures, described under "File Backup and Recovery" on page IG-94, to restore the file, or you can resume execution of your application. To restart your application, you can either IPL or cancel \$IAM and reload it.

If you wish to extend the logic of the error exit, code your own exit to replace the \$IAM task error exit. Then rename CDIERR (the \$IAM task error exit), name your error exit CDIERR, and rebuild \$IAM.

SYSTEM FUNCTION RETURN CODES

If a system function called by an Indexed Access Method request terminates with a positive return code, the return code is placed in a location named by the SYSRTCD parameter in the PROCESS or LOAD request. This location is used until a DISCONN is issued.

For example, the GET request uses the supervisor read function. If the read terminates with a positive return code, that return code is saved in the location named by the SYSRTCD parameter in the PROCESS request associated with the GET request. The GET request also terminates with a positive return code in the task control word. The positive return code indicates that a read error has occurred. The cause of the read error can be determined by examining the location named by the SYSRTCD parameter.

Note: When analyzing errors, the Indexed Access Method return code in the task code word would be checked prior to the system return code.

The following example is a method of obtaining the return code value from the location SYSRTCD. This routine gets the task SYSRTCD, and compares it to the EDX successful return code, negative one (-1).

```
YOURPRGM  PROGRAM  START
          .
          .
          .
          SUBROUT  ERRTEST
          MOVE     TASKRC,SYSRTCD      get system return code
          IF       (TASKRC,EQ,-1)      if -1, return now
          .
          .                               if not -1 then perform
          .                               your diagnosis
          ENDIF
          RETURN
          .
          .
TASKRC  DATA  F'0'                    saved system return cods
```

THE DATA-SET-SHUT-DOWN CONDITION

Sometimes an I/O error occurs that is not associated with a specific request. For example, task A issues a GET on file X. To secure buffer space to satisfy the request, the Indexed Access Method attempts to write a block to file Y and, in writing the record, an error occurs. Data set Y is damaged but there is no requesting program to accept an error return code.

The error is indicated by setting the data-set-shut-down condition for file Y. After this condition occurs, no requests except a DISCONN are accepted for file Y.

Later, if task B issues a GET on file Y, the request is terminated with a data-set-shut-down return code. Task B should issue a DISCONN and use recovery and backup procedures as described under "File Backup and Recov-

ery" on page IG-94, to reconstruct the file. To cancel the data-set-shut-down condition, initial program load (IPL) or cancel \$IAM.

\$ILOG - ERROR LOGGING FACILITY

You can have the Indexed Access Method to log errors in the system error log data set.

To use the error logging facility, allocate a system error log data set. The data set must be a minimum of 3 256-byte EDX records. The first two records are used for control, and this would allow one error log entry. The error log entries are entered in the log data set one after the other as they occur. When the data set becomes full, the new entries overlay the old entries starting at the front of the data set again. Therefore the size of the data set should be based on the frequency of errors, and the frequency with which the data set is listed or examined. Each error log entry requires a 256-byte EDX record.

You can load \$LOG into any partition. Using the system command LOGINIT, will initialize and activate error logging for any Indexed Access Method errors.

The following format is used to record Indexed Access Method associated errors in the error log data set:

Displacement into each EDX \$IAM error log entry

<u>Hex</u>	<u>(Dec)</u>	<u>Field</u>
00	(0)	Indexed Access Method entry identifier 'CDIIAM'
06	(6)	User TCB Address
08	(8)	Data Set Name
10	(16)	Volume
16	(22)	Original Function Byte
17	(23)	Current Function Byte
18	(24)	Indexed Access Method Return Code
1A	(26)	System Return Code
1C	(28)	User AKR Value

To list the Indexed Access Method error log entries currently in the system error log, you can load \$ILOG using the system command \$L. Respond to the prompt "(DSNAME,VOLUME):" with system error log data set and volume name.

The list will be directed to the terminal which was used to load \$ILOG.

Following is a sample of the printed error report showing two error records:

INDEXED ACCESS METHOD LOG REPORT PROGRAM ACTIVE										
TCB	PTN	ADDR	DSNAME	VOLUME	ORIG	CURR	SIAM	SYSTEM	DATE	TIME
					FNCTN	FNCTN	RTCODE	RTCODE		
2	1F64	IAMFILE	EDX002	PUTNW	PUTNW	62	-1	00/00/00	00:00:00	
2	1F64	IAMFILE	EDX002	PUTNW	PUTNW	62	-1	00/00/00	00:00:00	

2 INDEXED ACCESS METHOD LOG ENTRIES LOCATED

SILOG ENDED

DEADLOCKS AND THE LONG-LOCK-TIME CONDITION

Because the Indexed Access Method uses record and block locks to preserve file integrity, deadlock and long-lock-time conditions may occur.

The deadlock condition occurs when two or more tasks interact in such a way that one or more resources becomes permanently locked, making further progress impossible. A deadlock can also occur when two requests from the same task require a lock on the same record or a lock on the same block in sequential mode.

A long-lock-time condition occurs when your program acquires a record for update and does not return the record to \$IAM for a long time.

Application tasks should avoid using the Indexed Access Method in such a way that a record or block remains locked for a long period of time, because other tasks may attempt to use the same record or block. In a terminal-oriented system, make every effort to ensure that a record or block is not locked during operator "think" time. Specifically, you should attempt to follow these rules:

- Do not retrieve a record for update, display the record at the terminal, and wait for the operator to modify it.
- Do not retrieve a record in sequential mode, display the record at the terminal, and wait for an operator response.

In both of these cases, a record or block is locked during operator "think" time and could be locked indefinitely.

A deadlock cannot be broken except by freeing the locks (records) that are being waited on.

If your application uses more than one IACB, deadlocks are possible. For example, one task has read record A and attempts to read record B, while another task has read record B and attempts to read record A. If you are using more than one IACB per task, such as in Multiple Terminal Manager applications, use ENQ/DEQ and interprogram communications to avoid the deadlocks.

You can avoid the long-lock-time condition by using one of the following two methods:

1.
 - a. Retrieve the desired record without specifying update.
 - b. Perform processing in a work area.
 - c. Retrieve the record, specifying update.
 - d. Compare the record read in step 1 with the record read in step 3.

- e. If the records are identical, issue a PUTUP request, specifying the address of the copy in the work area. If they are not identical, issue a RELEASE request for the record read in step 3, and repeat steps 1 through 5.
2. Use conditional requests which do not wait for locks. (See "Chapter 7. Coding the Indexed Access Method Requests" for descriptions of coding conditional requests.)

To retrieve records in sequential mode, use the technique described in "Avoiding Resource Contention" on page IG-210.

VERIFYING REQUESTS AND FILES

Following are two steps you can take to help you isolate and correct malfunctions in your Indexed Access Method application program.

- Request verification - to determine that requests are correct check all parameters specified or defaulted on the Indexed Access Method CALL statements:
 - PROCESS/LOAD requests - When issuing a PROCESS or LOAD, check that the specified file name is the correct file control block (DSCB) for the file you are verifying.
 - GET-PUT-DELETE-RELEASE requests - For these requests, carefully check the key, its position, length, and the relational operator (if used). Ensure that the correct address for the indexed access control block (IACB) is passed from the PROCESS or LOAD request, and that the record area address is correct.
- File verification - read your \$VERIFY report or indexed file dump to determine whether data or index records are missing or incorrect.

Note: Be sure that the combination of parameters specified by the SE command of the \$IAMUT1 utility to define your file is correct (see the "Chapter 8. The \$IAMUT1 Utility" on page IG-141 for a description of the \$IAMUT1 parameters).



CHAPTER 12. INSTALLING THE INDEXED ACCESS METHOD

This chapter presents an overview of how to install the Indexed Access Method.

The Indexed Access Method is distributed on two double surface diskettes with external ID AM4001- and AM4002-V1M2BASVOL01. The diskettes are formatted are 256 bytes per sector.

INSTALLATION PROCEDURES

The installation information which follows is for planning purposes only. The specific details for installing the product is included in the "Program Directory", which is shipped with the product.

INSTALLING THE INDEXED ACCESS METHOD

Installing the Indexed Access Method consists of two steps:

1. Step 1
 - a. Ensure that adequate space is available for the installation according to the requirements shown in Figure 14.

Volume	Data Sets	EDX Records	Contents
EDX002	3	265	Load Modules
ASMLIB	2	145	Source Modules
ASMLIB	1	6	Link Module

Figure 14. Volume Space Requirements

2. Step 2
 - a. Copy the Indexed Access Method load module (\$IAM), the utility program (\$IAMUT1), the file verification program (\$VERIFY), and the log report program (\$ILOG), to the EDX002 volume.
 - b. Copy the following source modules to the ASMLIB volume.
 - IAMEQU
 - FCBEQU
 - IAM

ASSEMBLING AND EXECUTING THE INSTALLATION VERIFICATION PROGRAM

To assemble and execute the installation verification program:

1. Submit to the \$JOBUTIL utility, the 'proc' \$SAMPROC provided on volume AM4001 to assemble and link edit the verification program.

The source statements for the installation verification program are contained in a data set named SAMPLE on volume AM4001.

2. Use \$IAMUT1 to define and allocate an indexed file to be used by the installation verification program. Respond to the SE option 2 prompts with the indicated values:

BASEREC	10	FREEBLK	10
BLKSIZE	256	RSVBLK	0
RECSIZE	80	RSVIX	0
KEYSIZE	28	FPOOL	0
KEYPOS	1	DELTHR	0
FREEREC	1	DYN	10

3. Load \$SAMPLE and when prompted for the data set and volume, respond with the name for the file allocated in the previous step (SAMPFILE).

Note: The procedure \$SAMPROC assumes that ASMWORK and LINKWORK data sets exist on EDX002. Allocate these data sets if they don't already exist with the \$DISKUT1 Event Driven Executive Utility. (Refer to the Operator's Reference for details on allocation of these data sets.)

APPENDIX A. SUMMARY OF CALCULATIONS

The following calculations can be used to define an indexed data set. For a more detailed description of these calculations, see "Chapter 3. Defining Primary Index Files" on page IG-17. In the calculations requiring division, results with non-zero remainders are either:

truncated T , or rounded up U .

To truncate is to drop the remainder; to round up is to add one (only if the remainder is non-zero), and truncate.

Data Block

[1] Records per data block = block size minus 16, divided by record size; result truncated. T
[1] = (BLKSIZE-16)/RECSIZE

[2] Free records per block. [2] = FREEREC

[3] Allocated records per data block = Records per block minus free records per block. [3] = [1] - [2]

Index Block (General)

[4] Index entry size = key length plus 4; must be even; add 1 if odd. [4] = KEYSIZE + 4 (+1 if odd)

[5] Total entries per index block = block size 16, divided index entry size; result truncated. T
[5] = (BLKSIZE-16) / [4]

Index Block (PIXB)

[6] Free entries per primary index block (PIXB) = specified percentage of total entries per index block; result rounded up. U
[6] = FREEBLK % of [5]

[7] Reserve entries per PIXB = specified percentage of total entries per index block; result rounded up. If free entries per PIXB and reserve entries per all PIXB entries, subtract one from reserve entries per PIXB.

$$[7] = \text{RSVBLK} \% \text{ of } [5]$$

U

$$(-1 \text{ if } [6] + [7] = [5])$$

[8] Allocated entries per PIXB = total entries per index block minus free entries per PIXB, minus reserve entries per PIXB.

$$[8] = [5] - [6] - [7]$$

Index Block (SIXB)

[9] Reserve entries per secondary index block (SIXB) = specified percentage of total entries per index block; result rounded up. If reserve entries per SIXB require all SIXB entries, subtract one.

$$[9] = \text{RSVIX} \% \text{ of } [5]$$

U

$$(-1 \text{ if } [9] = [5])$$

[10] Allocated entries per SIXB = total entries per index block minus reserve entries per SIXB.

$$[10] = [5] - [9]$$

Delete Threshold

[11] The number of blocks to retain in cluster (delete threshold) is calculated in one of three ways:

- a. If the RSVBLK parameter was not specified:
Number of blocks to retain in cluster = total entries per index block.

$$[11] = [5]$$

or

- b. If the RSVBLK parameter was specified, but the DELTHR parameter was not specified: Number of blocks to retain in the cluster = allocated entries per PIXB, plus one-half of free entries per PIXB; result rounded up.

$$[11] = [8] + [6] / 2$$

U

or

c. If the RSVBLK parameter was specified, and the DELTHR parameter was specified: Number of blocks to retain in the cluster = specified percentage of total entries per index block; result rounded up. If the result is zero, set it to 1.

$$[11] = \text{DELTHR \% of } [5]$$

U

(If 0, set [11] to 1)

Data in Data Set

[12] Initial allocated data blocks = base records divided divided by allocated records records per data block; result rounded up.

$$[12] = \text{BASEREC} / [3]$$

U

[13] Number of clusters in data set = initial allocated data blocks, divided by allocated entries per PIXB; result rounded up.

$$[13] = [12] / [8]$$

U

[14] Total number of free blocks in data set = number of clusters in data set, times free entries per PIXB.

$$[14] = [13] \times [6]$$

Indexes in Data Set

[15] Number of primary index blocks (PIXBs) = number of clusters in data set.

$$[15] = [13]$$

[16]aNumber of secondary index blocks (SIXBs) = number of PIXBs, divided by allocated entries per SIXB; result rounded up.

$$[16] = [15] / [10]$$

U

[17]aCalculate the number of index blocks for levels 3 to n. Note that levels 1 (PIXB) and 2 (SIXB) have already been calculated. When the number of index blocks at a level is 1, n has been reached and the calculation is finished.

Number of index blocks at level i (i=3 to n) = number of index blocks at next lower level, divided by total entries per index block; result rounded up.

$$[17] = [17] / [5]$$

U

[18] Total number of index blocks = sum of index blocks at each level until a level containing a single index block is attained.

$$[18] = [15] + [16] + (\text{Sum of all } [17]\text{s})$$

Free Pool

[19] Number of new data blocks which can be assigned to existing clusters = reserve entries per PIXB, times number of PIXBs.

$$[19] = [7] \times [15]$$

[20] Number of new clusters (PIXBs) which can be created = reserve entries per SIXB, times number of SIXBs.

$$[20] = [9] \times [16]$$

[21] Number of new data blocks which can be assigned to new clusters = total entries per index block, times number of new clusters which can be created.

$$[21] = [5] \times [20]$$

[22] Maximum possible free pool = number of new data blocks data blocks which can be assigned to existing clusters, plus number of new clusters (PIXBs) which can be created, plus number of new data blocks which can be assigned to new clusters.

$$[22] = [19] + [20] + [21]$$

[23] Actual number of free pool blocks = specified percentage of maximum possible free pool; result rounded up.

$$[23] = \text{FPOOL \% of } [22]$$

U

Size of Data Set

[24] Total number of blocks in data set = 1 (for file control block), plus total number of index blocks, plus initial allocated data blocks, plus total number of free blocks in data set, plus actual number of free pool blocks.

$$\begin{aligned} [24] &= 1 + [18] + [12] \\ &+ [14] + [23] \end{aligned}$$



APPENDIX B. PREPARING INDEXED ACCESS METHOD PROGRAMS

To prepare an application program that issues Indexed Access Method requests, perform the following steps:

1. Enter your source program statements, using one of the Event Driven Executive text editors (\$FSEDIT, \$EDIT1, or \$EDIT1N).
2. Create the \$EDXLINK control statements required to combine your program with \$IAM (the link module) and any other object modules you may need in your application. Use one of the text editors to perform this operation.
3. Assemble or compile your source program.
4. Use the linkage editor, \$EDXLINK, to combine the object modules into a single load module, using the control statements prepared in Step 2.

When the preceding steps are completed, the program is ready to be executed.

Refer to the Program Preparation Guide for more information on program preparation.

A SAMPLE \$JOBUTIL PROCEDURE AND LINK-EDIT CONTROL DATA SET

The following are examples of a \$JOBUTIL procedure and a link-edit control file used to prepare a program.

Sample \$JOBUTIL Procedure

The following \$JOBUTIL procedure is an example of preparing an EDL program.

```
*****
*
* THESE STATEMENTS WILL COMPILE, LINK, AND UPDATE THE
* APPLICATION.
*
*****
JOB      COMPILE
***     COMPILE USERPROG SOURCE   ***
LOG      $SYSPRTR
PROGRAM  $EDXASM,ASMLIB
DS       USERPROG,EDX002          SOURCE MODULE
DS       ASMWORK,EDX002           ASSEMBLER WORK DATA SET
DS       USEROBJ,EDX002          ASSEMBLER OUTPUT
PARM     LIST      $SYSPRTR
EXEC
JUMP     END,GT,4
JOB      LINK
LOG      $SYSPRTR
PROGRAM  $EDXLINK,EDX002
DS       LINKWORK,EDX002          WORK DSNAME
*                                               LINK-CONTROL DATA SET
PARM     LINKCNTL,EDX002          &SYSPRTR
EXEC
LABEL   END
EOJ
```

Link Edit Control Data Set Example

The following link-edit control records can be used to link-edit an Indexed Access Method application with the Indexed Access Method.

```
*****  
*  
* LINK EDIT CONTROL DATA SET (LINKCTL)  
*  
*****  
INCLUDE USEROBJ,EDX002  INCLUDE APPLICATION PGM OBJECT  
INCLUDE IAM,ASMLIB      INCLUDE INDEXED ACCESS METHOD  
LINK    USERPROG,EDX002 REPLACE END  
END
```


APPENDIX C. CODING EXAMPLES

This chapter demonstrates how to code the Indexed Access Method request functions by means of sample programs. This example uses Event Driven Language CALL functions. The second example uses the COBOL language. The third example in this chapter is coded using PL/I language.

EDL INDEXED ACCESS METHOD CODING EXAMPLE

This program gives an example for each of the Indexed Access Method function calls. The indexed file is opened first in load mode and ten base records are loaded followed by a DISCONNECT. Next, the same file is opened for processing. A GET request is performed for the first record whose key is greater than 'JONES PW'. Two more records are retrieved sequentially and then the ENDSEQ call releases the file from sequential mode. A record is then retrieved directly by key and updated. Another record is retrieved sequentially and deleted. A new record is inserted and another one is deleted by their unique keys. Finally, an example of extracting information from the file control block is shown. Upon successful completion the message "Verification Complete" is displayed on the console.

Although using secondary keys are not demonstrated in this example the requests are coded the same for secondary keys as they are for primary keys. When accessing secondary keys use the secondary index file name instead of the primary index file name. The Indexed Access Method will open the primary index file and retrieve the data record according to the secondary key requested.

This program requires that an Indexed Access Method file has been defined with the \$IAMUT1 utility with the following specifications:

BASEREC	10
BLKSIZE	256
RECSIZE	80
KEYSIZE	28
KEYPOS	1
FREEREC	1
FREEBLK	10
RSVBLK	0
RSVIX	0
FPOOL	0
DELTHR	0
DYN	0

EDL INDEXED ACCESS METHOD CODING EXAMPLE

```
SAMPLE PROGRAM START,DS=??,ERRXIT=TEECB
START EQU *
*
ENQT
PRINTTEXT LOGON,LINE=0 PRINT LOGON MESSAGE
DEQT
*
* OPEN THE INDEXED ACCESS METHOD DATA SET FOR LOADING
*
CALL IAM,(LOAD),IACB,(DS1),(OPENTAB),(SHARE)
*
* LOAD THE INDEXED ACCESS METHOD DATA SET
*
MOVEA POINTER,RECORD1          POINTER <== A(RECORD1)
DO RECNUM,TIMES
CALL IAM,(PUT),IACB,(*),P4=POINTER
ADD POINTER,80                POINT TO NEXT RECORD
ENDDO
* GET OUT OF LOAD MODE
CALL IAM,(DISCONN),IACB
EJECT
*
* OPEN THE INDEXED FILE FOR PROCESSING
*
CALL IAM,(PROCESS),IACB,(DS1),(OPENTAB),(SHARE)
*
* PERFORM A DIRECT RETRIEVAL OF THE FIRST RECORD WHOSE KEY IS
* GREATER THAN 'JONES PW'. THE KEY FIELD WILL BE MODIFIED TO
* REFLECT THE KEY OF THE RECORD RETRIEVED.
*
CALL IAM,(GET),IACB,(BUFF),(KEY3),(GT)
MOVE RTCODE,SAMPLE
IF (SAMPLE,NE,-1),GOTO,IAMERR
*
* PERFORM A SEQUENTIAL RETRIEVAL OF THE FIRST TWO RECORDS
* WHOSE KEYS ARE GREATER THAN OR EQUAL TO 'JONES PW'
*
CALL IAM,(GETSEQ),IACB,(BUFF),(KEY1),(GE)
MOVE RTCODE,SAMPLE
IF (SAMPLE,NE,-1),GOTO,IAMERR
CALL IAM,(GETSEQ),IACB,(BUFF)
MOVE RTCODE,SAMPLE
IF (SAMPLE,NE,-1),GOTO,IAMERR
CALL IAM,(ENDSEQ),IACB,(BUFF) END SEQUENTIAL MODE
```

```

* UPDATE THE RECORD WHOSE KEY IS 'JONES PW' BY A DIRECT UPDATE
*
      CALL IAM,(GET),IACB,(BUFF),(KEY1),(UPEQ)
      MOVE RTCODE,SAMPLE
      IF (SAMPLE,NE,-1),GOTO,IAMERR
*
* MAKE THE DESIRED MODIFICATIONS TO THE RECORD NOW IN BUFFER
*
      MOVE BUFF+30,0
      CALL IAM,(PUTUP),IACB,(BUFF)
*
* DELETE THE RECORD WHOSE KEY IS 'JONES PW' BY A
* SEQUENTIAL UPDATE
*
      CALL IAM,(GETSEQ),IACB,(BUFF),(KEY1),(UPEQ)
      MOVE RTCODE,SAMPLE
      IF (SAMPLE,NE,-1),GOTO,IAMERR
      CALL IAM,(PUTDE),IACB,(BUFF)
      CALL IAM,(ENDSEQ),IACB          END SEQUENTIAL MODE
*
* INSERT A NEW RECORD WITH A KEY OF 'MATHIS GR'
*
      CALL IAM,(PUT),IACB,(NEWREC)
*
* DELETE THE RECORD WHOSE KEY IS 'LANG LK'
*
      CALL IAM,(DELETE),IACB,(KEY2)
      MOVE RTCODE,SAMPLE
      IF (SAMPLE,NE,-1),GOTO,IAMERR
      EJECT
*
* EXTRACT THE FILE CONTROL BLOCK INTO THE EXTRACT BUFFER
*
      CALL IAM,(EXTRACT),IACB,(EXTBUF),(FCBSIZE),128
      MOVEA #1,EXTBUF          #1 <-- A(EXTRACT BUFFER)
      MOVE FLAGBYTE,(0,#1),BYTE  OBTAIN FCB FLAG BYTE
      SPACE 5
*
* WRITE VERIFICATION COMPLETE MESSAGE TO THE OPERATOR
*
      ENQT
          PRINTTEXT SKIP=1
          PRINTTEXT VERIF,SPACES=0
      DEQT
      GOTO FINISH          JUMP AROUND ERROR ROUTINES
SYSERR EQU *          GETS CONTROL ON SYSIPGM CHECK

```



```

* WHEN A TASK ERROR EXIT IS SPECIFIED IN AN INDEXED
* ACCESS METHOD PROGRAM, YOU CAN RELEASE ALL ACTIVE
* RECORD AND BLOCK LEVEL LOCKS AS WELL AS DISCONNECT
* THE FILE ITSELF BY ISSUING THE 'DISCONN' CALL FOR
* EACH FILE THAT IS OPEN.

```

```

*
      GOTO FINISH
      EJECT
IAMERR EQU *           GETS CONTROL UPON INDEXED
*                               METHOD ERRORS

      MOVE RTCODE,SAMPLE
      ENQT

      PRINTTEXT SKIP=2
      PRINTTEXT RTCODMSG
      PRINTNUM RTCODE,TYPE=S,FORMAT=(3,0,I)
      PRINTTEXT SKIP=1
      PRINTTEXT ERRMSG,SPACES=0

      DEQT
FINISH EQU *
      CALL IAM,(DISCONN),IACB
      PROGSTOP
      EJECT

```

```

*
* DATA DEFINITION AND STORAGE AREAS
*

```

RECNUM	DATA	F'10'	NUMBER OF RECORDS TO LOAD
RTCODE	DATA	F'0'	INDEXED ACCESS METHOD RETURN CODE
OPENTAB	DATA	F'0'	SYSTEM RETURN CODE ADDRESS
	DATA	A(IAMERR)	ERROR EXIT ROUTINE ADDRESS
	DATA	F'0'	END OF DATA ROUTINE ADDRESS
RECORD1	DATA	CL80'BAKER RG'	
RECORD2	DATA	CL80'DAVIS EN'	
RECORD3	DATA	CL80'HARRIS SL'	
RECORD4	DATA	CL80'JONES PW'	
RECORD5	DATA	CL80'JONES TR'	
RECORD6	DATA	CL80'LANG LK'	
RECORD7	DATA	CL80'PORTER JS'	
RECORD8	DATA	CL80'SMITH AR'	
RECORD9	DATA	CL80'SMITH GA'	
RECORD10	DATA	CL80'THOMAS SN'	
FLAGBYTE	DATA	H'0'	FCB FLAG BYTE
	DATA	H'0'	

```

NEWREC  DATA  CL80'MATHIS GR'
BUFF    DATA  CL80' '
KEY1    TEXT   'JONES PW',LENGTH=28
KEY2    TEXT   'LANG LK',LENGTH=28
        DATA  X'1C'
        DATA  X'00'
KEY3    DATA  CL28'JONES PW'
IACB    DATA  F'0'          ADDR OF IACB PUT HERE
EXTBUF  DATA  64F'0'       FCB PUT HERE BY EXTRACT
LOGON   TEXT   'INSTALLATION VERFICATION PROGRAM ACTIVE'
VERIF   TEXT   'VERIFICATION COMPLETE'
ERRMSG  TEXT   'VERIFICATION INCOMPLETE DUE TO BAD RETURN CODES'
RTCODMSG TEXT  'INDEXED ACCESS METHOD RETURN CODE: '
        EJECT

```

```

*
* THE FOLLOWING STORAGE IS USED BY TASK ERROR EXIT HANDLING
*

```

```

TEECB   EQU    *           TASK ERROR EXIT CONTROL BLOCK
        DATA  F'2'       # OF DATA WORDS THAT FOLLOW
        DATA  A(SYSERR)  ADDRESS OF EXIT ROUTINE
        DATA  A(HSA)     ADDRESS OF HARDWARE STATUS AREA

```

```

*
* HARDWARE STATUS AREA. THIS STORAGE WILL BE FILLED IN BY
* HARDWARE UPON SYSTEM OR PROGRAM CHECK INTERRUPT
*

```

```

HSA     EQU    *           PROCESSOR STATUS WORD
        DATA  F'0'
HSALSB  EQU    *           LEVEL STATUS BLOCK:
        DATA  F'0'       ADDRESS KEY REGISTER
        DATA  F'0'       INSTRUCTION ADDR REGISTER
        DATA  F'0'       LEVEL STATUS REGISTER
        DATA  8F'0'      GENERAL REGISTERS 0-7
        COPY   IAMEQU
        COPY   FCBEQU
        ENDPROG
        END

```

COBOL INDEXED ACCESS METHOD CODING EXAMPLE

This coding example inserts, deletes, and updates records in an indexed file, using primary and secondary keys to retrieve the records. The indexed file description is described below under "Input File".

Program Description

This program reads a record and based on a transaction code, either updates, deletes, or inserts records to a current Indexed Access Method file. The transaction type also determines whether indexing is done using a secondary or primary key.

Input File

I. TRANSACTION FILE.

TRANSACTION RECORD FORMAT:

EMPLOYEE NUMBER	1-6	(6)
LAST NAME	7-21	(15)
FIRST NAME	22-31	(10)
ADDRESS	32-56	(25)
CITY	57-68	(12)
STATE	69-70	(2)
AGE	71-72	(2)
START DATE	73-78	(6)
TYPE	79	(1)
ACTION	80	(1)

II. UPDATE FILE

A. MASTER FILE.

PRIMARY. KEY IS EMPLOYEE NUMBER

B. NAME FILE

SECONDARY. KEY IS LAST NAME.

EMPLOYEE RECORD FORMAT

EMPLOYEE NUMBER	1-6	(6)
LAST NAME	7-21	(15)
FIRST NAME	22-31	(10)
ADDRESS	32-56	(25)
CITY	57-68	(12)
STATE	69-70	(2)
AGE	71-72	(2)
START DATE	73-78	(6)
FILLER	79-80	(2)

```

IDENTIFICATION DIVISION.
PROGRAM-ID.                COBOL1.
ENVIRONMENT DIVISION.
CONFIGURATION SECTION.
SOURCE-COMPUTER.          IBM-S1.
OBJECT-COMPUTER.          IBM-S1.
SPECIAL-NAMES.
    SYSOUT IS PRINTER.
INPUT-OUTPUT SECTION.
FILE-CONTROL.
    SELECT EMPLOYEE-MASTER ASSIGN TO DS2 "EMPMAS" "EDXTST"
        ORGANIZATION IS INDEXED
        ACCESS MODE IS RANDOM
        RECORD KEY IS EMPLOYEE-NUMBER
        FILE STATUS IS SK.
    SELECT EMP-NAME-FILE ASSIGN TO DS3 "EMPNAME" "EDXTST"
        ORGANIZATION IS INDEXED
        ACCESS MODE IS RANDOM
        RECORD KEY IS E-LAST-NAME
        FILE STATUS IS SK.
    SELECT TRANSACTION-FILE ASSIGN TO DS4 "TRANSF" "EDXTST"
        ORGANIZATION IS SEQUENTIAL
        ACCESS MODE IS SEQUENTIAL
        FILE STATUS IS SK.

DATA DIVISION.
FILE SECTION.
FD  EMPLOYEE-MASTER
    LABEL RECORDS ARE STANDARD
    BLOCK CONTAINS 3 RECORDS.
01  MASTER-RECORD.
    05  EMPLOYEE-NUMBER    PICTURE X(06).
    05  FILLER             PICTURE X(74).

FD  EMP-NAME-FILE
    LABEL RECORDS ARE STANDARD
    BLOCK CONTAINS 3 RECORDS.
01  EMP-NAME-RECORD.
    05  FILLER             PICTURE X(06).
    05  E-LAST-NAME       PICTURE X(15).
    05  FILLER             PICTURE X(59).

FD  TRANSACTION-FILE
    LABEL RECORDS ARE STANDARD
    BLOCK CONTAINS 3 RECORDS
    RECORD CONTAINS 80 CHARACTERS.
01  TRANS-ACTION-RECORD  PICTURE X(80).

```

WORKING-STORAGE SECTION.

77 EOF PICTURE 9(01) VALUE ZERO.
77 ERR-SWITCH PICTURE X(01) VALUE ZERO.
77 SK PICTURE X(02) VALUE ZERO.

01 EMPLOYEE-RECORD.

05 MAN-NUMBER PICTURE X(06).
05 NAME.
 10 LAST-NAME PICTURE X(15).
 10 FIRST-NAME PICTURE X(10).
05 STREET-ADDRESS PICTURE X(25).
05 CITY PICTURE X(12).
05 STATE PICTURE X(02).
05 AGE PICTURE X(02).
05 START-DATE PICTURE X(06).
05 FILLER PICTURE X(02).

01 TRANSACTION-RECORD.

05 T-NUMBER PICTURE X(06).
05 T-NAME.
 10 T-LAST-NAME PICTURE X(15).
 10 T-FIRST-NAME PICTURE X(10).
05 T-STREET-ADDRESS PICTURE X(25).
05 T-CITY PICTURE X(12).
05 T-STATE PICTURE X(02).
05 T-AGE PICTURE X(02).
05 T-START-DATE PICTURE X(06).
05 TRANSACTION-CODE.
 10 TRANS-TYPE PICTURE X(01).
 10 TRANS-ACTION PICTURE X(01).

PROCEDURE DIVISION.

BEGIN-PROCESSING.

PERFORM FILE-OPEN1.
IF SK = "00"
 PERFORM PROCESS-SECTION UNTIL EOF = 1.
PERFORM CLOSE-UP.
DISPLAY " CLOSE UP PROC COMPLETE" UPON PRINTER.
STOP RUN.

PROCESS-SECTION.

READ TRANSACTION-FILE INTO TRANSACTION-RECORD
AT END MOVE 1 TO EOF
DISPLAY "TRANSACTION FILE PROCESSING COMPLETE" UPON PRINTER.
IF EOF NOT = 1
 DISPLAY " TRANSACTION-RECORD UPON PRINTER
 PERFORM CONTROL-SECTION.

```

CONTROL-SECTION.
  IF TRANS-ACTION NOT = "2"
    PERFORM READ-SECTION.
  IF TRANS-ACTION = 1
    PERFORM DELETE-PROC
  ELSE
    IF TRANS-ACTION = 2
      PERFORM ADD-PROC
    ELSE
      PERFORM UPDATE-PROC.
READ-SECTION.
  IF TRANS-ACTION = "P"
    PERFORM READ-PRIMARY
  ELSE
    PERFORM READ-SECONDARY.

READ-PRIMARY.
  DISPLAY " DURING PRIMARY READ" UPON PRINTER.
  DISPLAY " KEY= " T-NUMBER UPON PRINTER.
  MOVE T-NUMBER TO EMPLOYEE-NUMBER.
  READ EMPLOYEE-MASTER INTO EMPLOYEE-RECORD
  INVALID KEY DISPLAY "INVALID PRIMARY KEY" T-NUMBER
  UPON PRINTER.
  IF SK NOT = "00"
    DISPLAY "PRIMARY READ FAILED " T-NUMBER " " SK UPON
    PRINTER.

READ-SECONDARY.
  DISPLAY " DURING SECONDARY READ" UPON PRINTER.
  DISPLAY " KEY= " T-LAST-NAME UPON PRINTER.
  MOVE T-LAST-NAME TO E-LAST-NAME.
  READ EMP-NAME-FILE INTO EMPLOYEE-RECORD
  INVALID KEY DISPLAY "INVALID SECONDARY KEY"
  T-LAST-NAME UPON PRINTER.
  IF SK NOT = "00"
    DISPLAY "SECONDARY READ FAILED" T-LAST-NAME " " SK
    UPON PRINTER.

UPDATE-PROC.
  IF T-LAST-NAME NOT = SPACES
    MOVE T-LAST-NAME TO LAST-NAME.
  IF T-FIRST-NAME NOT = SPACES
  IF T-STREET-ADDRESS NOT = SPACES
    MOVE T-STREET-ADDRESS TO STREET-ADDRESS.

```

```
IF T-CITY NOT = SPACES
    MOVE T-CITY TO CITY.
IF T-STATE NOT = SPACES
    MOVE T-STATE TO STATE.
IF T-AGE NOT = SPACES
    MOVE T-AGE TO AGE.
IF T-START-DATE NOT = SPACES
    MOVE T-START-DATE TO START-DATE.
IF TRANS-TYPE = "P"
    PERFORM PRIMARY-REWRITE
ELSE
    PERFORM SECONDARY-REWRITE.
```

PRIMARY-REWRITE.

```
    DISPLAY " BEGIN PRIMARY REWRITE KEY = " T-NUMBER
        UPON PRINTER.
    MOVE T-NUMBER TO EMPLOYEE-NUMBER.
    REWRITE MASTER-RECORD FROM EMPLOYEE-RECORD
        INVALID KEY DISPLAY "INVALID PRIMARY KEY"
        T-NUMBER UPON PRINTER.
    IF SK NOT = "00"
        DISPLAY "PRIMARY WRITE FAILED" T-NUMBER " " SK
        UPON PRINTER.
    DISPLAY " PRIMARY REWRITE COMPLETE" UOPN PRINTER.
```

SECONDARY-REWRITE.

```
    DISPLAY " BEGIN SECONDARY REWRITE KEY = " T-LAST-NAME
        UPON PRINTER.
    MOVE T-LAST-NAME TO E-LAST-NAME.
    REWRITE EMP-NAME-RECORD FROM EMPLOYEE-RECORD
        INVALID KEY DISPLAY "INVALID SECONDARY KEY"
        " " T-LAST-NAME UPON PRINTER.
    IF SK NOT = "00"
        DISPLAY "SECONDARY WRITE FAILED" T-LAST-NAME " "
        SK UPON PRINTER.
    DISPLAY " SECONDARY REWRITE COMPLETE" UPON PRINTER.
```

DELETE-PROC.

```
    IF TRANS-TYPE = "P"
        PERFORM PRIMARY-DELETE
    ELSE
        PERFORM SECONDARY-DELETE.
```

PRIMARY-DELETE.

MOVE TRANS-ACTION-RECORD TO MASTER-RECORD.
DISPLAY " PRIMARY DELETE STARTED" UPON PRINTER.
DELETE EMPLOYEE-MASTER RECORD
INVALID KEY DISPLAY "INVALID PRIMARY KEY"
T-NUMBER UPON PRINTER.
IF SK NOT = "00"
DISPLAY "PRIME KEY FOR DELETE NOT FOUND"
UPON PRINTER.
DISPLAY " PRIMARY DELETE FINISHED" UPON PRINTER.

SECONDARY-DELETE.

MOVE TRANS-ACTION-RECORD TO EMP-NAME-RECORD.
DISPLAY " SECONDARY DELETE STARTED" UPON PRINTER.
DELETE EMP-NAME-FILE RECORD
INVALID KEY DISPLAY "INVALID SECONDARY KEY"
T-LAST-NAME UPON PRINTER.
IF SK NOT = "00" DISPLAY
"SECONDARY KEY FOR DELETE NOT FOUND" UPON PRINTER
DISPLAY TRANSACTION-RECORD UPON PRINTER.
DISPLAY " SECONDARY DELETE FINISHED" UPON PRINTER.

ADD-PROC.

MOVE TRANSACTION-RECORD TO EMPLOYEE-RECORD
IF TRANS-TYPE = "P"
PERFORM PRIMARY-ADD
ELSE
PERFORM SECONDARY-ADD.

PRIMARY-ADD.

WRITE MASTER-RECORD FROM EMPLOYEE-RECORD
INVALID KEY DISPLAY "INVALID PRIMARY KEY"
T-NUMBER UPON PRINTER.
IF SK NOT = "00" DISPLAY
"INSERT FAILED FOR PRIME FILE" UPON PRINTER
DISPLAY TRANSACTION-RECORD UPON PRINTER.

SECONDARY-ADD.

WRITE EMP-NAME-RECORD FROM EMPLOYEE-RECORD
INVALID KEY DISPLAY "INVALID SECONDARY KEY"
T-LAST-NAME UPON PRINTER.
IF SK NOT = "00"
DISPLAY "INSERT FAILED FOR SECONDARY FILE"
UPON PRINTER.


```

FILE-OPEN1.
  OPEN I-O EMPLOYEE-MASTER.
  IF SK NOT = "00"
    DISPLAY "OPEN FAILED FOR EMPMAST" SK UPON PRINTER
  ELSE
    OPEN I-O EMP-NAME-FILE
    IF SK NOT = "00"
      DISPLAY "OPEN FAILED FOR EMPNAME" SK UPON PRINTER
    ELSE
      OPEN INPUT TRANSACTION-FILE
      IF SK NOT = "00"
        DISPLAY "OPEN FAILED FOR TRANSACTION-FILE" SK
          UPON PRINTER.
      DISPLAY " FILE OPEN COMPLETE" UPON PRINTER.

CLOSE-UP.
  DISPLAY " BEGIN CLOSE UP PROC " UPON PRINTER.
  CLOSE TRANSACTION-FILE.
  CLOSE EMP-NAME-FILE.
  IF SK NOT = "00"
    DISPLAY "CLOSE FAILED FOR EMPNAME, RC= " SK
      UPON PRINTER.
  ELSE
    DISPLAY " EMP-NAME-FILE CLOSED " UPON PRINTER.
  CLOSE EMPLOYEE-MASTER.
  IF SK NOT = "00"
    DISPLAY "CLOSE FAILED FOR EMPMAST, RC= " SK
      UPON PRINTER.
  ELSE
    DISPLAY " EMP-MAST-FILE CLOSED " UPON PRINTER.

```

PL/I INDEXED ACCESS METHOD CODING EXAMPLE

This PL/I coding example inserts, deletes, and updates records in an indexed file, using primary and secondary keys to retrieve the records. The indexed file description is described below under "Input File".

Program Description

This program reads a record and based on a transaction code, either updates, deletes, or inserts records to a current Indexed Access Method file. The transaction code also determines whether index access is done using a secondary key or primary key.

Input File.

I. INPUT FILE

A. TRANSACTION FILE,

TRANSACTION RECORD FORMAT:

EMPLOYEE NUMBER	1-6	(6)
LAST NAME	7-21	(15)
FIRST NAME	22-31	(10)
ADDRESS	32-56	(25)
CITY	57-68	(12)
STATE	69-70	(2)
AGE	71-72	(2)
START DATE	73-78	(6)
TYPE	79	(1)
ACTION	80	(1)

II. UPDATE FILE

A. MASTER FILE.

PRIMARY. KEY IS EMPLOYEE NUMBER

B. NAME FILE

SECONDARY. KEY IS LAST NAME.

EMPLOYEE RECORD FORMAT

EMPLOYEE NUMBER	1-6	(6)
LAST NAME	7-21	(15)
FIRST NAME	22-31	(10)
ADDRESS	32-56	(25)
CITY	57-68	(12)
STATE	69-70	(2)
AGE	71-72	(2)
START DATE	73-78	(6)
FILLER	79-80	(2)

```

PL1TEST: PROCEDURE OPTIONS(MAIN);
DCL EMPMAST                                /* EMPLOYEE MASTER FILE */
      FILE RECORD                          /* PRIMARY */
      DIRECT                               /* KEY IS EMPLOYEE NUMBER */
      UPDATE
      KEYED
      ENV(FB BLKSIZE(256) RECSIZE(80) INDEXED
           KEYLENGTH(6) KEYLOC(1));

DCL EMPNAME                                /* EMPLOYEE NAME FILE */
      FILE RECORD                          /* SECONDARY */
      DIRECT                               /* KEY IS EMPLOYEE NAME */
      UPDATE
      KEYED
      ENV(FB BLKSIZE(256) RECSIZE(80) INDEXED
           KEYLENGTH(15) KEYLOC(7));

DCL TRANSFL                                /* TRANSACTION FILE */
      FILE RECORD                          /* INPUT FILE */
      SEQUENTIAL
      INPUT
      ENV(FB BLKSIZE(240) RECSIZE(80) CONSECUTIVE);

DCL SYSPRINT                               /* STANDARD OUTPUT FILE */
      FILE PRINT
      ENV(F BLKSIZE(121));

DCL 1 $FCBLST                              /* FILE CONTROL BLOCK LIST */
      STATIC
      EXTERNAL,
      2 $FCBCNT                             /* FILE COUNT */
      FIXED BIN(15)
      INIT(4),
      2 $FCBF1                              /* FILE #1 NAME */
      CHAR(8)
      INIT('EMPMAST'),
      2 $FCBD1                              /* FILE #1 DATA SET */
      CHAR(8)
      INIT('EMPMAST'),
      2 $FCBV1                              /* FILE #1 VOLUME */
      CHAR(6)
      INIT('EDXTST'),
      2 $FCBF2                              /* FILE #2 NAME */
      CHAR(8)
      INIT('EMPNAME'),
      2 $FCBD2                              /* FILE #2 DATA SET */
      CHAR(8)
      INIT('EMPNAME'),

```

```

2 $FCBV2                /* FILE #2 VOLUME      */
    CHAR(6)
    INIT('EDXTST'),
2 $FCBF3                /* FILE #3 NAME        */
    CHAR(8)
    INIT('TRANSFL'),
2 $FCBD3                /* FILE #3 DATA SET   */
    CHAR(8)
    INIT('TRANSFL'),
2 $FCBV3                /* FILE #3 VOLUME      */
    CHAR(6)
    INIT('EDXTST'),
2 $FCBF4                /* FILE #4 NAME        */
    CHAR(8)
    INIT('SYSPRINT'),
2 $FCBD4                /* FILE #4 DEVICE NAME */
    CHAR(8)
    INIT('SYSPRINT'),
2 $FCBTL4              /* FILE #4 TOP LINE    */
    FIXED BIN(15)
    INIT(1),
2 $FCBBL4              /* FILE #4 BOTTOM LINE */
    FIXED BIN(15)
    INIT(66),
2 $FCBHL4              /* NOT USED            */
    FIXED BIN(15)
    INIT(00);

DCL 1 EMP_RECORD       /* IAM BUFFER FORMAT   */
    STATIC,
    2 EMP_NUMBER
        CHAR(6),
    2 NAME,
        3 LAST_NAME
            CHAR(15),
        3 FIRST_NAME
            CHAR(10),
    2 STREET_ADDRESS
        CHAR(25),
    2 CITY
        CHAR(12),
    2 STATE
        CHAR(2),
    2 AGE
        CHAR(2),
    2 START_DATE
        CHAR(6),
    2 FILLER
        CHAR(2);

```

```

DCL 1 TRANSACTIONR                               /* TRANSACTION RECORD FORMAT */
  STATIC,
  2 TEMP_NUMBER
    CHAR(6),
  2 TNAME,
  3 TLAST_NAME
    CHAR(15),
  3 TFIRST_NAME
    CHAR(10),
  2 TSTREET_ADDRESS
    CHAR(25),
  2 TCITY
    CHAR(12),
  2 TSTATE
    CHAR(2),
  2 TAGE
    CHAR(2),
  2 TSTART_DATE
    CHAR(6),
  2 TRANSACTION_CODE,
  3 TRANS_TYPE                                /* 'P' = PRIMARY 'S' = SECONDARY */
    CHAR(1),
  3 TRANS_ACTION                                /* 1 = DELETE 2 = INSERT */
    CHAR(1);                                /* 3 = UPDATE */
DCL BLANK                                        /* FIELD OF BLANKS */
  STATIC
  CHAR (6)
  INIT(' ');

DCL (IOERR,
FOERR)                                          /* INPUT/OUTPUT ERROR FLAG */
  STATIC                                        /* OPEN ERROR FLAG */
  CHAR(1)
  INIT('F');

DCL TRUE                                        /* TRUE VALUE FOR FLAGS */
  STATIC
  CHAR(1)
  INIT('T'),
FALSE                                          /* FALSE VALUE FOR FLAGS */
  STATIC
  CHAR(1)
  INIT('F');

DCL R_CODE                                      /* RETURN CODE */
  STATIC
  FIXED BIN(15)
  INIT(0);

```

```

DCL ONCODE                /* ON CONDITION CODE          */
    BUILTIN;
DCL EOF                    /* END OF FILE FLAG          */
    STATIC
    CHAR(1)
    INIT('F');
/***** MAIN PROGRAM *****/
/*
/***** ON CONDITION FOR EOF *****/
    ON ENDFILE(TRANSFL)
    BEGIN;
        PUT LIST('*** TRANSACTION FILE PROCESSING COMPLETE ***');
        EOF = TRUE;
        CLOSE                /* CLOSE ALL FILES          */
            FILE(EMPMAST),
            FILE(EMPNAME),
            FILE(TRANSFL);
        STOP TASK;
    END;
/*
/***** ON CONDITIONS FOR FILE OPEN ERRORS *****/
    ON UNDF(EMPMAST) FOERR = TRUE;
    ON UNDF(EMPNAME) FOERR = TRUE;
    ON UNDF(TRANSFL) FOERR = TRUE;
/*
/***** ON CONDITIONS FOR I/O ERRORS *****/
    ON KEY(EMPMAST) IOERR = TRUE;
    ON KEY(EMPNAME) IOERR = TRUE;
/*
/***** OPEN ALL FILES *****/
/*
    CALL OPEN;
/***** INITIATE PROCESSING UNTIL EOF CONDITION IS REACHED *****/
DO WHILE (EOF ^= TRUE);
    IOERR = FALSE;
    FOERR = FALSE;
    CALL PROCESS;          /* INVOKE PROCESS SUBROUTINE */
END; /* END DO WHILE */
/***** END MAIN PROGRAM *****/
/*
OPEN: PROC;
    OPEN FILE(EMPMAST) UPDATE;
    IF FOERR = TRUE THEN
        DO;
            R_CODE = ONCODE;          /* SET RETURN CODE */
            PUT LIST('OPEN FAILED FOR EMPMAST') SKIP;
            PUT LIST('ON CODE = ',R_CODE) SKIP;
        END;

```

```

IF FOERR = FALSE THEN
DO;
  OPEN FILE(EMPNAME) UPDATE;
  IF FOERR = TRUE THEN
  DO;
    R_CODE = ONCODE;          /* SET RETURN CODE */
    PUT LIST('OPEN FAILED FOR EMPNAME') SKIP;
    PUT LIST('ON CODE = ',R_CODE) SKIP;
  END;
END;
/*
IF FOERR = FALSE THEN
DO;
  OPEN FILE(TRANSFL);
  IF FOERR = TRUE THEN
  DO;
    R_CODE = ONCODE;          /* SET RETURN CODE */
    PUT LIST('OPEN FAILED FOR TRANSFL') SKIP;
    PUT LIST('ON CODE = ',R_CODE);
  END;
END;
/*
IF FOERR = TRUE
THEN
  STOP TASK;
/*
END; /* END OPEN PROCEDURE */
/***** PROCESS PROCEDURE *****/
/*
/* 1) READS IN A TRANSACTION RECORD
/* 2) IF ACTION = 1 DELETES RECORD WITH CORRESPONDING KEY.
/*    = 2 INSERTS RECORD ONTO IAM FILE.
/*    = 3 READS RECORD WITH CORRESPONDING KEY,
/*      ALLOWS UPDATE, REWRITES RECORD.
/* 3) IF TYPE = 'P' ALL INDEXING IS DONE WITH A PRIMARY KEY
/*    = 'S' ALL INDEXING IS DONE WITH A SECONDARY KEY
/* 4) ALL IDENTIFIERS, FILES AND RECORDS USED ARE GLOBAL
/*
/*****
PROCESS: PROCEDURE;
  READ FILE(TRANSFL) INTO (TRANSACTIONR);
  IF IOERR = TRUE THEN
  DO;
    R_CODE = ONCODE;
    PUT LIST('READ HAS FAILED FOR TRANSFL') SKIP;
    PUT LIST('ON CODE = ',R_CODE) SKIP;
  END;

```

```

IF TRANS_ACTION = '1' & IOERR = FALSE THEN
    CALL DELETE;                               /* BEGIN DELETE */
IF TRANS_ACTION = '3' & IOERR = FALSE THEN
    CALL UPDATE;
IF TRANS_ACTION = '2' & IOERR = FALSE THEN
    CALL INSERT;

END; /* END PROCEDURE PROCESS */

DELETE: PROC;
    IF TRANS_TYPE = 'P'
        THEN
            DELETE FILE(EMPMAST) KEY(TEMP_NUMBER);
        ELSE
            DELETE FILE(EMPNAME) KEY(TLAST_NAME);
    END; /* END DELETE */
/*
UPDATE: PROC;
    IF TRANS_TYPE = 'P'
        THEN
            CALL PRIM_READ;
        ELSE
            CALL SEC_READ;
    IF IOERR = FALSE THEN
        DO;
            IF TLAST_NAME ^= BLANK
                THEN
                    LAST_NAME = TLAST_NAME;
            IF TFIRST_NAME ^= BLANK
                THEN
                    FIRST_NAME = TFIRST_NAME;
            IF TSTREET_ADDRESS ^= BLANK
                THEN
                    STREET_ADDRESS = TSTREET_ADDRESS;
            IF TCITY ^= BLANK
                THEN
                    CITY = TCITY;
            IF TSTATE ^= BLANK
                THEN
                    STATE = TSTATE;

```



```

        IF TAGE ^= BLANK
            THEN
                AGE = TAGE;
        IF TSTART_DATE ^= BLANK
            THEN
                START_DATE = TSTART_DATE;
        CALL REWRITE;

    END;                                /* END UPDATE          */
END;
/*                                     */
PRIM_READ: PROC;
    READ FILE(EMPMAST) INTO(EMP_RECORD)
        KEY(TEMP_NUMBER);
    IF IOERR = TRUE THEN
        DO;
            R_CODE = ONCODE;
            PUT LIST('EMPMAST PRIMARY READ HAS FAILED') SKIP;
            PUT LIST('KEY = ',TEMP_NUMBER);
            PUT LIST('ONCODE = ',R_CODE);
        END;
    END;                                /* END PRIMARY READ    */
/*                                     */
SEC_READ: PROC;
    READ FILE(EMPNAME) INTO(EMP_RECORD)
        KEY(TLAST_NAME);
    IF IOERR = TRUE THEN
        DO;
            R_CODE = ONCODE;
            PUT LIST('EMPMAST SECONDARY READ HAS FAILED') SKIP;
            PUT LIST('KEY = ',TLAST_NAME);
            PUT LIST('ONCODE = ',R_CODE);
        END;
    END;                                /* END SECONDARY READ  */
REWRITE: PROC;
    IF TRANS_TYPE = 'P'                 /* BEGIN REWRITE       */
        THEN
            REWRITE FILE(EMPMAST) FROM(EMP_RECORD)
                KEY(TEMP_NUMBER);
        ELSE
            REWRITE FILE(EMPNAME) FROM(EMP_RECORD)
                KEY(TLAST_NAME);
    END;

```

```

INSERT: PROC;
  IF TRANS_TYPE = 'P'
  THEN
    DO;
      WRITE FILE(EMPMAS) FROM(TRANSACTIONR)
        KEYFROM(TEMP_NUMBER);
      IF IOERR = TRUE THEN
        DO;
          R_CODE = ONCODE;
          PUT LIST('EMPMAS SECONDARY INSERTION HAS FAILED')
            SKIP;
          PUT LIST('KEY = ',TEMP_NUMBER);
          PUT LIST('ONCODE = ',R_CODE);
        END;
    END;
  ELSE
    DO;
      WRITE FILE(EMPNAME) FROM(TRANSACTIONR)
        KEYFROM(TLAST_NAME);
      IF IOERR = TRUE THEN
        DO;
          R_CODE = ONCODE;
          PUT LIST('EMPMAS SECONDARY INSERTION HAS FAILED')
            SKIP;
          PUT LIST('KEY = ',TNAME);
          PUT LIST('ONCODE = ',R_CODE);
        END;
    END;
  END;
END PL1TEST;

```



BIBLIOGRAPHY

This bibliography lists and summarizes the Event Driven Executive Version 3 library; it should assist you in using the library as a whole as well as direct you to the individual books you require.

See the publication IBM Series/1 Publications Directory, GC34-0361 for a complete list of Series/1 hardware and programming manuals. This publication lists each version of the Event Driven Executive and its related licensed programs, and shows which manuals apply to each version.

EVENT DRIVEN EXECUTIVE LIBRARY

The IBM Series/1 Event Driven Executive library consists of nine full-sized books, six specialized internal books, and a quick reference pocket book:

- IBM Series/1 Event Driven Executive System Guide, SC34-1702 (or System Guide)
- IBM Series/1 Event Driven Executive Operator's Reference, SC34-1703 (or Operator's Reference)
- IBM Series/1 Event Driven Executive Language Reference, SC34-1706 (or Language Reference)
- IBM Series/1 Event Driven Executive Communications and Terminal Applications Guide, SC34-1705 (or Communications and Terminal Applications Guide)
- IBM Series/1 Event Driven Executive Commercial Applications Development Guide, SC34-0381 (or Commercial Applications Development Guide)
- IBM Series/1 Event Driven Executive Program Preparation Guide, SC34-1704 (or Program Preparation Guide)
- IBM Series/1 Event Driven Executive Messages and Codes, SC34-0403 (or Messages and Codes)
- IBM Series/1 Event Driven Executive Systems Network Architecture and Remote Job Entry Guide, SC34-0402 (or Systems Network Architecture and Remote Job Entry Guide)
- IBM Series/1 Event Driven Executive Indexed Access Method Version 2 Guide, SC34-0404 (or Indexed Access Method Version 2 Guide)
- IBM Series/1 Event Driven Executive Internal Design, LY34-0212 (or Internal Design)

- IBM Series/1 Event Driven Executive System/370 Channel Attach Program Internal Design, LY34-0215 (or System/370 Channel Attach Program Internal Design),
- IBM Series/1 Event Driven Executive Multiple Terminal Manager Internal Design, LY34-0190 (or Multiple Terminal Manager Internal Design),
- IBM Series/1 Event Driven Executive Indexed Access Method Internal Design, LY34-0189 (or Indexed Access Method Internal Design),
- IBM Series/1 Event Driven Executive Indexed Access Method Version 2 Internal Design, LY24-0238 (or Indexed Access Method Version 2 Internal Design)
- IBM Series/1 Event Driven Executive Systems Network Architecture Internal Design, LY34-0615 (or Systems Network Architecture Internal Design)
- IBM Series/1 Event Driven Executive Reference Summary, SX34-0038 (or Reference Summary),

SUMMARY OF LIBRARY

System Guide

The System Guide has four parts: Concepts and Facilities, System Generation, Problem Determination, and Data Management.

The concepts and facilities part describes the Event Driven Executive system and its related products and introduces the capabilities of these products. The system generation part presents step-by-step procedures for generating a supervisor tailored to your Series/1 hardware configuration and programming needs. The problem determination part explains how to use Event Driven Executive debugging aids. Data management explains direct access data sets, tape data sets, partitioned data sets, dynamic data set allocation, and sensor I/O.

Operator's Reference

Operator's Reference describes:

- Event Driven Executive utility programs
- Operator commands
- Procedures to execute system and application programs
- The session manager -- a menu-driven interface program that will invoke the programs required for program development

The operator commands and utilities are presented in alphabetic order and include detailed syntax and explanations. Also included is a description of the facilities of the session manager.

Utilities are described either in detail or with a reference to the publication where a more detailed description can be found.

Language Reference

The Language Reference familiarizes you with the Event Driven Language by first grouping the instructions into functional categories. Then the instructions are presented alphabetically, with complete syntax and an explanation of each operand. Instructions that support special functions such as 370 Channel Attach are documented in the manual that describes the function. An appendix summarizes the syntax of all instructions.

The final section of the Language Reference contains examples of using the Event Driven Language to code functions such as:

- Program loading
- User exit routine

- I/O level control program
- Indexing and hardware register usage

Communications and Terminal Applications Guide

The Communications and Terminal Applications Guide introduces the Event Driven Executive communications and terminal support.

The Communications and Terminal Applications Guide describes:

- Binary synchronous communications
- Asynchronous communications
- The Host Communications Facility
- Series/1 to Series/1 Attachment support
- The Remote Management Utility
- The System/370 Channel Attach Program
- General Purpose Interface Bus support.

The Communications and Terminal Applications Guide contains the utilities and coding details on Event Driven Language instructions needed for communications support and advanced terminal applications such as graphics.

Commercial Applications Development Guide

The Commercial Applications Development Guide describes:

- The Indexed Access Method Version 1
- The Multiple Terminal Manager
- Printer spooling
- High Level Language interface to static screen subroutines

The description of the Indexed Access Method explains how to plan, design, and build indexed data sets. Also included is information on how to write applications that use indexed data sets.

The Multiple Terminal Manager section explains how to write applications to run under the Multiple Terminal Manager and how to use the Multiple Terminal Manager.

The description of printer spooling explains how to install and use the spooling facility, and describes the spool utility and spool operator commands.

Program Preparation Guide

The Program Preparation Guide describes:

- The Event Driven Language compiler
- The Series/1 assembler
- The host assembler
- The linkage editor
- The object program converter (native and host)
- The data set and overlay prefind utility

SNA and Remote Job Entry Guide

Systems Network Architecture and Remote Job Entry Guide describes the Systems Network Architecture services and functions that are available through a high-level instruction set, and explains Systems Network Architecture formats and protocols. The manual also describes how to use the Systems Network Architecture Remote Job Entry utility.

Messages and Codes

Messages and Codes describes many of the informational, warning, and error messages produced by the Event Driven Executive system and its related products. The manual also contains return codes, completion codes, and post codes issued by these products.

Indexed Access Method Version 2 Guide

This manual describes Version 2 of the Indexed Access Method. It explains how to set up and access an indexed file, and how to use the Indexed Access Method utilities. The manual also contains examples of using the Indexed Access Method in Event Driven Language, COBOL, and PL/I programs.

Internal Designs

Internal Design describes the internal logic flow and specifications of the Event Driven Executive system so that you can understand how the system interfaces with application programs. It familiarizes you with the design and implementation by describing the purpose, function, and operation of the various Event Driven Executive system programs.

Multiple Terminal Manager Internal Design, System/370 Channel Attach Internal Design, Indexed Access Method Internal Design, Indexed Access Method Version 2 Internal Design, and Systems Network Architecture Inter-

nal Design describe the internal logic flow and specifications of these programs.

Unlike the other manuals in the library, the Internal Design books contain material that is the licensed property of IBM and they are available only to licensed users of the products that they describe.

Reference Summary

The Reference Summary is a pocket-sized booklet to be used for quick reference. It lists the Event Driven Language instructions with their syntax, the utility and program preparation commands, and completion, post, and return codes.

OTHER EVENT DRIVEN EXECUTIVE PROGRAMMING PUBLICATIONS

The following is a list of Event Driven Executive publications that you may need to code application programs in languages other than the Event Driven Language.

- IBM Series/1 FORTRAN IV Language Reference, GC34-0133.
- IBM Series/1 Event Driven Executive FORTRAN IV User's Guide, SC34-0315.
- IBM Series/1 Mathematical and Functional Subroutine Library User's Guide, SC34-0139.
- IBM Series/1 Event Driven Executive PL/I Language Reference, GC34-0147.
- IBM Series/1 Event Driven Executive PL/I User's Guide, SC34-0148.
- IBM Series/1 COBOL Language Reference, GC34-0234.
- IBM Series/1 COBOL Version 2 Language Reference, SC34-0392
- IBM Series/1 Event Driven Executive COBOL Programmer's Guide, SL23-0014.
- IBM Series/1 Event Driven Executive COBOL Version 2 Programmer's Guide, SC34-0393
- IBM Series/1 Event Driven Executive Macro Assembler Reference, GC34-0317.
- IBM Series/1 Macro Assembler Reference Summary, SX34-0128

OTHER PROGRAMMING PUBLICATIONS

The following publications may be required for reference.

- IBM Series/1 Programming System Summary, GC34-0285.
- IBM Series/1 Event Driven Executive Study Guide, SR30-0436.
- IBM Data Processing Glossary, GC20-1699.
- DOS/VS Tape Labels Manual, GC33-5374
- IBM Series/1 Event Driven Executive Sort/Merge Programmer's Guide, SL23-0016
- IBM Series/1 Data Collection Interactive Programming RPQ P82600 User's Guide, SC34-1654.
- IBM Series/1 Host Communications Facility Program Description Manual, SH20-1819.
- IBM OS/VS Basic Telecommunications Access Method (BTAM), GC27-6980.
- General Information — Binary Synchronous Communications, GA27-3004.
- IBM System/370 Program Preparation Facility, SB30-1072.

SERIES/1 SYSTEM LIBRARY PUBLICATIONS

The following publications contain hardware-related information.

- IBM Series/1 4952 Processor and Processor Features Description, GA34-0084.
- IBM Series/1 4953 Processor and Processor Features Description, GA34-0022.
- IBM Series/1 4955 Processor and Processor Features Description, GA34-0021.
- IBM Series/1 Communications Features Description, GA34-0028.
- IBM 3101 Display Terminal Description, GA18-2033.
- IBM Series/1 4962 Disk Storage Unit and 4964 Diskette Unit Description, GA34-0024.
- IBM Series/1 4963 Disk Subsystem Description, GA34-0051.
- IBM Series/1 4966 Diskette Magazine Unit Description, GA34-0052.
- IBM Series/1 4969 Magnetic Tape Subsystem Description, GA34-0087.

- IBM Series/1 4973 Line Printer Description, GA34-0044.
- IBM Series/1 4974 Printer Description, GA34-0025.
- IBM Series/1 4978-1 Display Station (RPQ D02055) and Attachment (RPQ D02038) General Information, GA34-1550
- IBM Series/1 4978-1 Display Station, Keyboard (RPQ D02056) General Information, GA34-1551
- IBM Series/1 4978-1 Display Station, Keyboard (RPQ D02057) General Information, GA34-1552
- IBM Series/1 4978-1 Display Station Keyboards (RPQ D02064 and D02065) General Information, GA34-1553
- IBM Series/1 4979 Display Station Description, GA34-0026
- IBM Series/1 4982 Sensor Input/Output Unit Description, GA34-0027
- IBM Diskette - General Information Manual, GA21-9182
- IBM Series/1 General Purpose Interface Bus (GPIB) Adapter - RPQ D02118, GA34-1556
- IBM Series/1 to Series/1 Attachment RPQs D02241 & D02242 Custom Feature, GA34-1561
- IBM Series/1 System/370 Channel Attachment Feature and 4993 Model 1 Series/1 System/370 Termination Enclosure Description, GA34-0057
- IBM Series/1 Data Collection Interactive RPQs D02312, D02313, and D02314 Custom Feature, GA34-1567

GLOSSARY OF TERMS AND ABBREVIATIONS

This glossary contains terms and abbreviations used in the Series/1 Event Driven Executive software publications. All software and hardware terms are Series/1 oriented. This glossary also serves as a supplement to the IBM Data Processing Glossary (GC20-1699).

\$\$SYSLOGA, \$\$SYSLOGB: The name of the alternate system logging device. This device is optional but, if defined, should be a terminal with keyboard capability, not just a printer.

\$\$SYSLOG: The name of the system logging device or operator station; must be defined for every system. It should be a terminal with keyboard capability, not just a printer.

\$\$SYSPRTR: The name of the system printer.

ACCA: See asynchronous communications control adapter.

address key: Identifies a set of Series/1 segmentation registers and represents an address space. It is one less than the partition number.

address space: The logical storage identified by an address key. An address space is the storage for a partition.

application program manager: The component of the Multiple Terminal Manager that provides the program management facilities required to process user requests. It controls the contents of a program area and the execution of programs within the area.

application program stub: A collection of subroutines that are

appended to a program by the linkage editor to provide the link from the application program to the Multiple Terminal Manager facilities.

asynchronous communications control adapter: An ASCII terminal attached via #1610, #2091 with #2092, or #2095 with #2096 adapters.

attention list: A series of pairs of 1 to 8 byte EBCDIC strings and addresses pointing to EDL instructions. When the attention key is pressed on the terminal, the operator can enter one of the strings to cause the associated EDL instructions to be executed.

backup: A copy of data to be used in the event the original data is lost or damaged.

base record slots: Space in an indexed file that is reserved for based records to be placed.

base records: Records are placed into an indexed file while in load mode or inserted in process mode with a new high key.

basic exchange format: A standard format for exchanging data on diskettes between systems or devices.

binary synchronous device data block (BSCDDB): A control block that provides the information to control one Series/1 Binary Synchronous Adapter. It determines

the line characteristics and provides dedicated storage for that line.

block: (1) See data block or index block. (2) In the Indexed Method, the unit of space used by the access method to contain indexes and data.

block mode: The transmission mode in which the 3101 Display Station transmits a data stream, which has been edited and stored, when the SEND key is pressed.

BSCDDB: See binary synchronous device data block.

buffer: An area of storage that is temporarily reserved for use in performing an input/output operation, into which data is read or from which data is written. See input buffer and output buffer.

bypass label processing: Access of a tape without any label processing support.

CCB: See terminal control block.

central buffer: The buffer used by the Indexed Access Method for all transfers of information between main storage and indexed files.

character image: An alphabetic, numeric, or special character defined for an IBM 4978 Display Station. Each character image is defined by a dot matrix that is coded into eight bytes.

character image table: An area containing the 256 character images that can be defined for an IBM 4978 Display Station. Each character image is coded into eight bytes, the entire table of codes requiring 2048 bytes of storage.

character mode: The transmission mode in which the 3101 Display Sta-

tion immediately sends a character when a keyboard key is pressed.

cluster: In an indexed file, a group of data blocks that is pointed to from the same primary-level index block, and includes the primary-level index block. The data records and blocks contained in a cluster are logically contiguous, but are not necessarily physically contiguous.

COD (change of direction): A character used with ACCA terminal to indicate a reverse in the direction of data movement.

cold start: Starting the spool facility by erasing any spooled jobs remaining in the spool data set from any previous spool session.

command: A character string from a source external to the system that represents a request for action by the system.

common area: A user-defined data area that is mapped into every partition at the same address. It can be used to contain control blocks or data that will be accessed by more than one program.

completion code: An indicator that reflects the status of the execution of a program. The completion code is displayed or printed on the program's output device.

controller: A device that has the capability of configuring the GPIB bus by designating which devices are active, which devices are listeners, and which device is the talker. In Series/1 GPIB implementation, the Series/1 is always the controller.

conversion: See update.

cross partition service: A function that accesses data in two partitions.

data block: In an indexed file, an area that contains control information and data records. These blocks are a multiple of 256 bytes.

data record: In an indexed file, the records containing customer data.

data set: A group of records within a volume pointed to by a directory member entry in the directory for the volume.

data set control block (DSCB): A control block that provides the information required to access a data set, volume or directory using READ and WRITE.

data set shut down: An indexed data set that has been marked (in main storage only) as unusable due to an error.

DCE: See directory control entry.

DDB: See disk data block.

direct access: (1) The access method used to READ or WRITE records on a disk or diskette device by specifying their location relative the beginning of the data set or volume. (2) In the Indexed Access Method, locating any record via its key without respect to the previous operation. (3) A condition in terminal I/O where a READTEXT or a PRINTEXT is directed to a buffer which was previously enqueued upon by an IOCB.

directory: (1) A series of contiguous records in a volume that describe the contents in terms of allocated data sets and free spaces. (2) For the Indexed Access Method Version 2, a data set that defines the relationship between

primary and secondary indexed files (secondary index support).

directory control entry (DCE): The first 32 bytes of the first record of a directory in which a description of the directory is stored.

directory member entry (DME): A 32-byte directory entry describing an allocated data set or volume.

disk data block (DDB): A control block that describes a direct access volume.

display station: An IBM 4978, 4979, or 3101 display terminal or similar terminal with a keyboard and a video display.

DME: See directory member entry.

DSCB: See data set control block.

dynamic storage: An increment of storage that is appended to a program when it is loaded.

end-of-data indicator: A code that signals that the last record of a data set has been read or written. End-of-data is determined by an end-of-data pointer in the DME or by the physical end of the data set.

ECB: See event control block.

EDL: See Event Driven Language.

emulator: The portion of the Event Driven Executive supervisor that interprets EDL instructions and performs the function specified by each EDL statement.

end-of-tape (EOT): A reflective marker placed near the end of a tape and sensed during output. The marker signals that the tape is nearly full.

event control block (ECB): A control block used to record the status (occurred or not occurred) of an event; often used to synchronize the execution of tasks. ECBs are used in conjunction with the WAIT and POST instructions.

event driven language (EDL): The language for input to the Event Driven Executive compiler (\$EDXASM), or the Macro and Host assemblers in conjunction with the Event Driven Executive macro libraries. The output is interpreted by the Event Driven Executive emulator.

EXIO (execute input or output): An EDL facility that provides user controlled access to Series/1 input/output devices.

external label: A label attached to the outside of a tape that identifies the tape visually. It usually contains items of identification such as file name and number, creation data, number of volumes, department number, and so on.

external name (EXTRN): The 1- to 8-character symbolic EBCDIC name for an entry point or data field that is not defined within the module that references the name.

FCA: See file control area.

FCB: See file control block.

file: A set of related records treated as a logical unit. Although file is often used interchangeably with data set, it usually refers to an indexed or a sequential data set.

file control area (FCA): A Multiple Terminal Manager data area that describes a file access request.

file control block (FCB): The first block of an indexed file. It contains descriptive information about the data contained in the file.

file control block extension: The second block of an indexed file. It contains the file definition parameters used to define the file.

file manager: A collection of sub-routines contained within the program manager of the Multiple Terminal Manager that provides common support for all disk data transfer operations as needed for transaction-oriented application programs. It supports indexed and direct files under the control of a single callable function.

formatted screen image: A collection of display elements or display groups (such as operator prompts and field input names and areas) that are presented together at one time on a display device.

free pool: In an indexed data set, a group of blocks that can be used for either data blocks or index blocks. These differ from other free blocks in that these are not initially assigned to specific logical positions in the file.

free space: In an indexed file, records blocks that do not currently contain data, and are available for use.

free space entry (FSE): An 8-byte directory entry defining an area of free space within a volume or a device.

FSE: See free space entry.

general purpose interface bus: The IEEE Standard 488-1975 that allows various interconnected devices to be attached to the GPIB adapter (RPQ D02118).

GPB: See general purpose interface bus.

group: A unit of 100 records in the spool data set allocated to a spool job.

H exchange format: A standard format for exchanging data on diskettes between systems or devices.

hardware timer: The timer features available with the Series/1 processors. Specifically, the 7840 Timer Feature card or the native timer (4952 only). Only one or the other is supported by the Event Driven Executive.

host assembler: The assembler licensed program that executes in a 370 (host) system and produces object output for the Series/1. The source input to the host assembler is coded in Event Driven Language or Series/1 assembler language. The host assembler refers to the System/370 Program Preparation Facility (5798-NNQ).

host system: Any system whose resources are used to perform services such as program preparation for a Series/1. It can be connected to a Series/1 by a communications link.

IACB: See indexed access control block.

IAR: See instruction address register.

ICB: See indexed access control block.

IIB: See interrupt information byte.

image store: The area in a 4978 that contains the character image table.

immediate data: A self-defining term used as the operand of an instruction. It consists of numbers, messages or values which are processed directly by the computer and which do not serve as addresses or pointers to other data in storage.

index: In an indexed file, an ordered collection of pairs of keys and pointers, used to sequence and locate records.

index block: In an indexed file, an area that contains control information and index entries. These blocks are a multiple of 256 bytes.

indexed access control block (IACB/ICB): The control block that relates an application program to an indexed file.

indexed access method: An access method for direct or sequential processing of fixed-length records by use of a record's key.

indexed data set: Synonym for indexed file.

indexed file: A file specifically created, formatted and used by the Indexed Access Method. An indexed file is sometimes called an indexed data set.

index entry: In an indexed file, a key-pointer pair, where the pointer is used to locate a lower-level index block or a data block.

index register (#1, #2): Two words defined in EDL and contained in the task control block for each task. They are used to contain data or for address computation.

input buffer: (1) See buffer. (2) In the Multiple Terminal Manager, an area for terminal input and output.

input output control block (IOCB): A control block containing information about a terminal such as the symbolic name, size and shape of screen, the size of the forms in a printer, or an optional reference to a user provided buffer.

instruction address register (IAR): The pointer that identifies the instruction currently being executed. The Series/1 maintains a hardware IAR to determine the Series/1 assembler instruction being executed. It is located in the level status block (LSB).

interactive: The mode in which a program conducts a continuous dialogue between the user and the system.

internal label: An area on tape used to record identifying information (similar to the identifying information placed on an external label). Internal labels are checked by the system to ensure that the correct volume is mounted.

interrupt information byte (IIB): In the Multiple Terminal Manager, a word containing the status of a previous input/output request to or from a terminal.

invoke: To load and activate a program, utility, procedure, or subroutine.

job: A collection of related program execution requests presented in the form of job control statements, identified to the jobstream processor by a JOB statement.

job control statement: A statement in a job that specifies requests for program execution, program parameters, data set definitions, sequence of execution, and, in general, describes the

environment required to execute the program.

job stream processor: The job processing facility that reads job control statements and processes the requests made by these statements. The Event Driven Executive job stream processor is \$JOBUTIL.

jumper: A wire or pair of wires which are used for the arbitrary connection between two circuits or pins in an attachment card.

key: In the Indexed Access Method, one or more consecutive characters used to identify a record and establish its order with respect to other records. See also key field. See also key field.

key field: A field, located in the same position in each record of an indexed file, whose content is used for the key of a record.

level status block (LSB): A Series/1 hardware data area that contains processor status.

library: A set of contiguous records within a volume. It contains a directory, data sets and/or available space.

line: A string of characters accepted by the system as a single input from a terminal; for example, all characters entered before the carriage return on the teletypewriter or the ENTER key on the display station is pressed.

link edit: The process of resolving external symbols in one or more object modules. A link edit can be performed with \$EDXLINK, whose output is a loadable program, or with \$LINK, whose output is a single, composite object module that is the input to the update process.

listener: A controller or active device on a GPIB bus that is configured to accept information from the bus.

load mode: In the Indexed Access Method, the mode in which records are loaded into base record slots in an indexed file.

load module: A single module having cross references resolved and prepared for loading into storage for execution. The module is the output of the \$UPDATE or \$UPDATEH utility.

load point: (1) Address in the partition where a program is loaded. (2) A reflective marker placed near the beginning of a tape to indicate where the first record is written.

lock: In the Indexed Access Method, a method of indicating that a record or block is in use and is not available for another request.

logical screen: A screen defined by margin settings, such as the TOPM, BOTM, LEFTM and RIGHTM parameters of the TERMINAL or IOCB statement.

LSB: See level status block.

member: A term used to identify a named portion of a partitioned data set (PDS). Sometimes member is also used as a synonym for a data set. See data set.

menu: A formatted screen image containing a list of options. The user selects an option to invoke a program.

menu-driven: The mode of processing in which input consists of the responses to prompting from an option menu.

multifile volume: A unit of recording media, such as tape reel or disk pack, that contains more than one data file.

multiple terminal manager: An Event Driven Executive licensed program that provides support for transaction-oriented applications on a Series/1. It provides the capability to define transactions and manage the programs that support those transactions. It also manages multiple terminals as needed to support these transactions.

multivolume file: A data file that, due to its size, requires more than one unit of recording media (such as tape reel or disk pack) to contain the entire file.

new high key: A key higher than any other key in an indexed file.

non-labeled tapes: Tapes that do not contain identifying labels (as in standard labeled tapes) and contain only files separated by tape-marks.

null character: A user-defined character used to define the unprotected fields of a formatted screen.

option selection menu: A full screen display used by the Session Manager to point to other menus or system functions, one of which is to be selected by the operator. (See primary option menu and secondary option menu.)

output buffer: (1) See buffer. (2) In the Multiple Terminal Manager, an area used for screen output and to pass data to subsequent transaction programs.

overlay: The technique of reusing a single storage area allocated to a program during execution. The

storage area can be reused by loading it with overlay programs that have been specified in the PROGRAM statement of the program or by calling overlay segments that have been specified in the OVERLAY statement of \$EDXLINK.

overlay area: A storage area within a program reserved for overlay programs specified in the PROGRAM statement or overlay segments specified in the OVERLAY statement in \$EDXLINK.

overlay program: A program in which certain control sections can use the same storage location at different times during execution. An overlay program can execute concurrently as an asynchronous task with other programs and is specified in the EDL PROGRAM statement in the main program.

overlay segment: A self-contained portion of a program that is called and sequentially executes as a synchronous task. The entire program that calls the overlay segment need not be maintained in storage while the overlay segment is executing. An overlay segment is specified in the OVERLAY statement of \$EDXLINK.

parameter selection menu: A full screen display used by the Session Manager to indicate the parameters to be passed to a program.

partition: A contiguous fixed-sized area of storage. Each partition is a separate address space.

physical timer: Synonym for hardware timer.

post code: An indicator returned to the event control block to indicate the completion of a wait or an event.

prefind: To locate the data sets or overlay programs to be used by a program and to store the necessary information so that the time required to load the prefounded items is reduced.

primary file: An indexed file containing the data records and primary index.

primary file entry: For the Indexed Access Method Version 2, an entry in the directory describing a primary file.

primary index: The index portion of a primary file. This is used to access data records when the primary key is specified.

primary key: In an indexed file, the key used to uniquely identify a data record.

primary-level index block: In an indexed file, the lowest level index block. It contains the relative block numbers (RBNs) and high keys of several data blocks. See cluster.

primary menu: The program selection screen displayed by the Multiple Terminal Manager.

primary option menu: The first full screen display provided by the Session Manager.

primary task: The first task executed by the supervisor when a program is loaded into storage. It is identified by the PROGRAM statement.

priority: A combination of hardware interrupt level priority and a software ranking within a level. Both primary and secondary tasks will execute asynchronously within the system according to the priority assigned to them.

process mode: In the Indexed Access Method, the mode in which records can be retrieved, updated, inserted or deleted.

processor status word (PSW): A 16-bit register used to (1) record error or exception conditions that may prevent further processing and (2) hold certain flags that aid in error recovery.

program: A disk- or diskette-resident collection of one or more tasks defined by a PROGRAM statement; the unit that is loaded into storage. (See primary task and secondary task.)

program header: The control block found at the beginning of a program that identifies the primary task, data sets, storage requirements and other resources required by a program.

program/storage manager: A component of the Multiple Terminal Manager that controls the execution and flow of application programs within a single program area and contains the support needed to allow multiple operations and sharing of the program area.

protected field: A field in which the operator cannot use the keyboard to enter, modify, or erase data.

PSW: See processor status word.

QCB: See queue control block.

QD: See queue descriptor.

QE: See queue element.

queue control block (QCB): A data area used to serialize access to resources that cannot be shared. See serially reusable resource.

queue descriptor (QD): A control block describing a queue built by the DEFINEQ instruction.

queue element (QE): An entry in the queue defined by the queue descriptor.

record: (1) The smallest unit of direct access storage that can be accessed by an application program on a disk or diskette using READ and WRITE. Records are 256 bytes in length. (2) In the Indexed Access Method, the logical unit that is transferred between \$IAM and the user's buffer. The length of the buffer is defined by the user.

recovery: The use of backup data to recreate data that has been lost or damaged.

reflective marker: A small adhesive marker attached to the reverse (nonrecording) surface of a reel of magnetic tape. Normally, two reflective markers are used on each reel of tape. One indicates the beginning of the recording area on the tape (load point), and the other indicates the proximity to the end of the recording area (EOT) on the reel.

relative record number: An integer value identifying the position of a record in a data set relative to the beginning of the data set. The first record of a data set is record one, the second is record two, the third is record three.

reorganize: The process of copying the data in an indexed file to another indexed file in a manner that rearranges the data for more optimum processing and free space distribution.

restart: Starting the spool facility by processing spool jobs

remaining in the spool data set from any previous spool session.

return code: An indicator that reflects the results of the execution of an instruction or subroutine. The return code is placed in the task code word (at the beginning of the task control block).

roll screen: A display screen which is logically segmented into an optional history area and a work area. Output directed to the screen starts display at the beginning of the work area and continues on down in a line-by-line sequence. When the work area gets full, the operator presses ENTER/SEND and its contents are shifted into the optional history area and the work area itself is erased. Output now starts again at the beginning of the work area.

SBIOCB: See sensor based I/O control block.

second-level index block: In an indexed data set, the second-lowest level index block. It contains the addresses and high keys of several primary-level index blocks.

secondary file.: See secondary index.

secondary index: For the Indexed Access Method Version 2, an indexed file used to access data records by their secondary keys. Sometimes called a secondary file.

secondary index entry: For the Indexed Access Method Version 2, this an an entry in the directory describing a secondary index.

secondary key: For the Indexed Access Method Version 2, the key used to uniquely identify a data record.

secondary option menu: In the Session Manager, the second in a series of predefined procedures grouped together in a hierarchical structure of menus. Secondary option menus provide a breakdown of the functions available under the session manager as specified on the primary option menu.

secondary task: Any task other than the primary task. A secondary task must be attached by a primary task or another secondary task.

sector: The smallest addressable unit of storage on a disk or diskette. A sector on a 4962 or 4963 disk is equivalent to an Event Driven Executive record. On a 4964 or 4966 diskette, two sectors are equivalent to an Event Driven Executive record.

sensor based I/O control block (SBIOCB): A control block containing information related to sensor I/O operations.

sequential access: The processing of a data set in order of occurrence of the records in the data set. (1) In the Indexed Access Method, the processing of records in ascending collating sequence order of the keys. (2) When using READ/WRITE, the processing of records in ascending relative record number sequence.

serially reusable resource (SRR): A resource that can only be accessed by one task at a time. Serially reusable resources are usually managed via (1) a QCB and ENQ/DEQ statements or (2) an ECB and WAIT/POST statements.

service request: A device generated signal used to inform the GPIB controller that service is required by the issuing device.

session manager: A series of pre-defined procedures grouped together as a hierarchical structure of menus from which you select the utility functions, program preparation facilities, and language processors needed to prepare and execute application programs. The menus consist of a primary option menu that displays functional groupings and secondary option menus that display a breakdown of these functional groupings.

shared resource: A resource that can be used by more than one task at the same time.

shut down: See data set shut down.

source module/program: A collection of instructions and statements that constitute the input to a compiler or assembler. Statements may be created or modified using one of the text editing facilities.

spool job: The set of print records generated by a program (including any overlays) while engaged to a printer designated as a spool device.

spool session: An invocation and termination of the spool facility.

SRQ: See service request.

standard labels: Fixed length 80-character records on tape containing specific fields of information (a volume label identifying the tape volume, a header label preceding the data records, and a trailer label following the data records).

static screen: A display screen formatted with predetermined protected and unprotected areas. Areas defined as operator prompts or input field names are protected

to prevent accidental overlay by input data. Areas defined as input areas are not protected and are usually filled in by an operator. The entire screen is treated as a page of information.

subroutine: A sequence of instructions that may be accessed from one or more points in a program.

supervisor: The component of the Event Driven Executive capable of controlling execution of both system and application programs.

system configuration: The process of defining devices and features attached to the Series/1.

SYSGEN: See system generation.

system generation: The processing of user selected options to create a supervisor tailored to the needs of a specific Series/1 configuration.

system partition: The partition that contains the supervisor (partition number 1, address space 0).

talker: A controller or active device on a GPIB bus that is configured to be the source of information (the sender) on the bus.

tapemark: A control character recorded on tape used to separate files.

task: The basic executable unit of work for the supervisor. Each task is assigned its own priority and processor time is allocated according to this priority. Tasks run independently of each other and compete for the system resources. The first task of a program is the primary task. All tasks attached by the primary task are secondary tasks.

task code word: The first two words (32 bits) of a task's TCB; used by the emulator to pass information from system to task regarding the outcome of various operations, such as event completion or arithmetic operations.

task control block (TCB): A control block that contains information for a task. The information consists of pointers, save areas, work areas, and indicators required by the supervisor for controlling execution of a task.

task supervisor: The portion of the Event Driven Executive that manages the dispatching and switching of tasks.

TCB: See task control block.

terminal: A physical device defined to the EDX system using the TERMINAL configuration statement. EDX terminals include directly attached IBM displays, printers and devices that communicate with the Series/1 in an asynchronous manner.

terminal control block (CCB): A control block that defines the device characteristics, provides temporary storage, and contains links to other system control blocks for a particular terminal.

terminal environment block (TEB): A control block that contains information on a terminal's attributes and the program manager operating under the Multiple Terminal Manager. It is used for processing requests between the terminal servers and the program manager.

terminal screen manager: The component of the Multiple Terminal Manager that controls the presentation of screens and communications between terminals and transaction programs.

terminal server: A group of programs that perform all the input/output and interrupt handling functions for terminal devices under control of the Multiple Terminal Manager.

terminal support: The support provided by EDX to manage and control terminals. See terminal.

trace range: A specified number of instruction addresses within which the flow of execution can be traced.

transaction oriented applications: Program execution driven by operator actions, such as responses to prompts from the system. Specifically, applications executed under control of the Multiple Terminal Manager.

transaction program: See transaction-oriented applications.

transaction selection menu: A Multiple Terminal Manager display screen (menu) offering the user a choice of functions, such as reading from a data file, displaying data on a terminal, or waiting for a response. Based upon the choice of option, the application program performs the requested processing operation.

unprotected field: A field in which the operator can use the keyboard to enter, modify or erase data. Also called non-protected field.

update: (1) To alter the contents of storage or a data set. (2) To convert object modules, produced as the output of an assembly or compilation, or the output of the linkage editor, into a form that can be loaded into storage for program execution and to update the

directory of the volume on which the loadable program is stored.

user exit: (1) Assembly language instructions included as part of an EDL program and invoked via the USER instruction. (2) A point in an IBM-supplied program where a user written routine can be given control.

vary offline: (1) To change the status of a device from online to offline. When a device is offline, no data set can be accessed on that

device. (2) To place a disk or diskette in a state where it is unknown by the system.

vary online: To place a device in a state where it is available for use by the system.

volume: A disk, diskette, or tape subdivision defined using \$INITDSK or \$TAPEUT1.

volume label: A label that uniquely identifies a single unit of storage media.



COMMON INDEX

This index is common to the Event Driven Executive library. It is a combination of entries from the publications listed below. Each entry references all of the publications that contain information about the entry, so this index provides a cross-reference between the publications.

Each page number entry contains a single or double letter prefix which identifies the publication where the subject can be found. The letter prefixes have the following meanings:

- AM = Indexed Access Method Internal Design (Version 1)
- CA = Commercial Applications Development Guide
- CT = Communications and Terminal Application Guide
- ID = Internal Design
- IG = Indexed Access Method Version 2 Guide
- IM = Indexed Access Method Version 2 Internal Design
- LR = Language Reference
- M = Multiple Terminal Manager Internal Design
- MC = Messages and Codes
- OR = Operator's Reference
- PP = Program Preparation Guide
- SG = System Guide
- SI = SNA Internal Design
- SM = Sort/Merge Programmer's Guide
- SN = SNA and Remote Job Entry Guide
- T = S/370 Channel Attach Internal Design

Special Characters

\$\$EDXIT task error exit routine
 description SG-253
 extending the routine SG-255
 output example SG-254
 using SG-254
\$\$EDXLIB system name SG-284
\$\$EDXVOL system name LR-268,
 SG-284
\$A operator command OR-24, SG-61
\$AUTO data set PP-87, SG-420
\$B operator command OR-26, SG-61
\$BSTRCE utility CT-235, OR-90,
 SG-63
\$BSCUT1 utility CT-237, OR-91,
 SG-63
\$BSCUT2 utility CT-239, OR-92,
 SG-63
\$C operator command OR-27, SG-61
\$CAPGM program CT-382
\$CHANUT1 utility
 commands
 CA change device
 address CT-409
 EN terminate utility CT-409
 PR print trace area CT-409
 SP stop channel attach
 device CT-410
 ST start channel attach
 device CT-410
 TR enable/disable
 trace CT-410
 description CT-408, OR-94
 function SG-63
 invoking CT-408
\$COBOL utility SG-66
\$COMPRES utility OR-95, SG-63
\$COPY utility OR-100, SG-63
\$COPYUT1 utility OR-111, SG-63
\$CP operator command
 function SG-61
 syntax OR-29
\$D operator command OR-30, SG-61
\$DASDI utility OR-119, SG-63
\$DEBUG utility
 description OR-142
 with PASSTHRU (\$RMU) CT-314
\$DICOMP display composer
 command descriptions CT-165
 create partitioned data set
 member SG-322
 function SG-65
 invoking CT-165
 overview OR-165
\$DIINTR utility CT-186, OR-167,
 SG-65
\$DISKUT1 utility
 command description OR-168
 create partitioned data
 set SG-322
 function SG-63
 used with \$MIGAID
 utility SG-441
\$DISKUT2 utility
 commands OR-181
 description OR-180
 function SG-63
 printing I/O error logs SG-258
\$DISKUT3 program
 description SG-339
 function SG-57
 input to SG-340
 request block contents SG-341
 return codes MC-225, SG-343
 Version 3 conversion consider-
 ations SG-435
\$DIUTIL utility CT-192, OR-191,
 SG-322
\$DUMP utility OR-193, SG-65
\$E operator command
 function SG-61
 with spool writer
 active CA-407
\$EDIT1/\$EDITIN text editors
 control keys OR-202
 data set requirements OR-199
 function SG-64
 line editing commands OR-235
 overview OR-199
\$EDXASM Event Driven Language
 compiler
 completion codes MC-208, PP-21
 control statements ID-243
 description PP-11
 example SG-414
 function SG-66
 internals ID-233
 invoking
 with \$JOBUTIL PP-19
 with \$L PP-15
 with session manager PP-18
 language control data
 set PP-12
 listing program
 (\$EDXLIST) PP-24
 multiple overlay areas ID-249
 options PP-13
 output PP-20
 overlay programs ID-246
 overview PP-11, SG-71
 required data sets PP-11
 syntax checking ID-246
 use considerations
 \$COPY definitions PP-22
 instructions requiring
 support modules PP-23
 using the compiler PP-13
 work data set ID-275
\$EDXATSR supervisor routine ID-39
\$EDXDEF hardware configuration
 editing SG-180
 internals ID-8
 storage map ID-8
\$EDXL language control data
 set PP-12
 command syntax
\$EDXLINK AUTOCALL control state-
 ment PP-51
\$EDXLINK comment control
 statement PP-51
\$EDXLINK completion codes MC-209,
 PP-73
\$EDXLINK control statements PP-49
\$EDXLINK COPY control
 statement PP-52
\$EDXLINK data sets PP-48
\$EDXLINK END control
 statement PP-53
\$EDXLINK INCLUDE control
 statement PP-53
\$EDXLINK LINK control
 statement PP-54
\$EDXLINK linkage editor
 \$JOBUTIL interface PP-65
 \$L interactive interface PP-62

\$L non-interactive interface PP-61
 advantages SG-74
 AUTOCALL option
 data set PP-74
 data set record format PP-74
 processing PP-75
 system data set PP-75
 completion codes MC-209
 control statements
 AUTOCALL PP-51
 comment PP-51
 COPY PP-52
 END PP-53
 INCLUDE PP-53
 LINK PP-54
 OVERLAY PP-55
 RESET PP-56
 VOLUME PP-56
 data sets PP-48
 description OR-241
 function SG-66
 interactive \$JOBUTIL interface PP-67
 interactive session manager interface PP-68
 internals
 ESD item types ID-300
 hash table ID-304
 include table ID-306
 modules ID-311
 OCT table entry ID-309
 invoke using
 \$JOBUTIL PP-65
 \$JOBUTIL interactive PP-67
 \$JOBUTIL non-interactive PP-66
 \$L interactive PP-62
 \$L non-interactive PP-61
 session manager PP-68
 link map PP-70
 non-interactive \$JOBUTIL interface PP-66
 non-interactive session manager interface PP-68
 operator termination PP-69
 overview SG-72, SG-73
 primary control statement data set
 description PP-48
 example of PP-57
 requesting a link edit PP-60
 secondary control statement data set
 description PP-48
 example of PP-59
 specifying dynamic storage PP-77
 using \$LINK control statement data sets PP-76
 \$EDXLINK OVERLAY control statement PP-55
 \$EDXLINK RESET control statement PP-56
 \$EDXLINK VOLUME control statement PP-56
 \$EDXLIST program PP-24, SG-66
 \$EDXNUC supervisor data set
 in link-edit PP-89
 in system generation SG-192
 \$FONT utility SG-66
 \$FORT utility SG-66
 \$FSEDIT full-screen editor
 data set requirements OR-260
 function SG-64
 options
 BROWSE OR-265
 EDIT OR-266
 END OR-272
 READ OR-268
 SUBMIT OR-270
 WRITE OR-269
 options and commands
 summary OR-263
 overview OR-259
 primary commands OR-268
 program function (PF) keys OR-262
 scrolling OR-261
 \$GPIBUT1 utility
 commands CT-134
 description CT-134
 example CT-145
 overview SG-66
 use in debugging applications CT-144
 \$HCFUT1 utility CT-367, OR-294, SG-64
 \$HXUT1 utility SG-63
 \$IAMSTGM module removal IG-206
 \$IAMUT1 utility
 commands
 AL IG-151
 BF IG-144
 CR CA-105
 DE IG-153
 DF CA-106, IG-146
 DI CA-108, IG-149
 DR IG-150
 EC CA-109, IG-160
 EF IG-161
 IE IG-155
 LE IG-156
 LO CA-110, IG-162
 NP IG-166
 PG IG-167
 PP IG-168
 PS IG-169
 RE CA-113, IG-170
 RO CA-112, IG-171
 SE CA-114, IG-173
 UE IG-158
 UN CA-120, IG-184
 completion codes CA-121, IG-186, MC-210
 description CA-103
 function SG-63
 overview CA-14, OR-317
 \$ID statement
 coding description LR-36
 for problem determination SG-239
 \$IDEF \$EDXASM instruction definition
 description ID-267
 instruction format ID-252
 \$ILOG IG-216
 \$IMAGE subroutines
 example of using CT-55
 general description CT-44
 use for device independence in terminal support CT-41
 \$IMAGE utility
 commands OR-325
 description OR-321
 example SG-404
 function SG-66

program function keys OR-340
 \$IMDATA subroutine
 description CT-53
 example CT-42, CT-54, CT-55,
 SG-391
 return codes CT-54, MC-239
 \$IMDEFN subroutine
 description CT-48
 example CT-49, CT-55, SG-391
 \$IMOPEN subroutine
 description CT-46
 example CT-41, CT-47, CT-55,
 SG-390
 return codes CT-47, MC-239
 \$IMPROT subroutine
 description CT-50
 example CT-42, CT-51, CT-55,
 SG-391
 return codes CT-52, MC-240
 \$INDEX subroutine, \$EDXASM ID-259
 \$INITDSK utility
 commands OR-345
 description OR-343
 function SG-63
 used in Version 3
 conversion SG-438
 \$INITIAL dataset
 description SG-174
 with session manager OR-46
 \$IOTEST utility OR-368, SG-65
 \$JOBUTIL job stream processor
 commands OR-377
 completion codes MC-212
 example SG-424
 setup procedure OR-376
 usage example OR-396
 \$JOBUTIL utility
 using to invoke
 Sort/Merge SM-42
 \$L operator command
 description OR-32
 function SG-61
 syntax OR-32
 using to invoke
 sort/merge SM-40
 \$LEMSG data set PP-87
 \$LINK
 autocall facility PP-101
 autocall option PP-87
 combining program modules PP-5
 completion codes MC-213, PP-97
 control records PP-88
 data set requirements PP-86
 description OR-398, PP-85
 elimination of duplicate
 CSECTs PP-6
 example SG-419
 formatting modules for
 \$UPDATE PP-85
 function SG-66
 invoking
 with \$JOBUTIL PP-95
 with \$L PP-93
 with session manager PP-94
 multiple CSECTs PP-6
 object module record
 format ID-381
 output PP-95
 output listing PP-96
 overview PP-5, SG-72
 storage map PP-96
 use considerations PP-4
 \$LOADER program
 description ID-16

 location SG-206
 module description ID-72
 \$LOG utility
 description ID-393, OR-400
 function SG-65
 log data set OR-400
 \$LPARSE subroutine,
 \$EDXASM ID-266
 \$MIGRID utility
 commands SG-444
 description SG-441
 error handling SG-452
 examples SG-445
 overview SG-436
 \$MIGCOPY utility
 description SG-458
 examples SG-460
 overview SG-436
 \$MIGRATE utility
 description SG-457
 overview SG-436
 \$MOVEVOL utility OR-402, SG-63
 \$MTMUT1 utility
 description OR-409
 function SG-66
 \$NETCMD (application link to
 SNA) SN-12
 \$P operator command OR-33, SG-61
 \$PACK subroutine CT-59
 \$PDS utility
 allocating a data set SG-322
 command descriptions SG-334
 example SG-333
 return codes MC-226
 \$PFMAP utility OR-410, SG-66
 \$PL/I utility SG-66
 \$PREFIND utility
 description PP-121
 function SG-66
 use considerations PP-129
 \$PRT2780 utility CT-343, OR-413,
 SG-64
 \$PRT3780 utility CT-343, OR-414
 \$RECON utility OR-415, SG-66
 \$RJESNA utility
 adding a error exit
 routine SN-217
 control record
 processing SN-206
 error handling SN-216
 hardware and software require-
 ments SN-202
 host considerations SN-225
 how to install SN-221
 how to start SN-207
 how to terminate SN-216
 job entry systems
 supported SN-202
 messages SN-239
 operator commands SN-210
 overview SG-64, SN-201
 receiving compacted
 data SN-223
 sample sessions SN-218
 task error exit SN-217
 workstation features SN-203
 workstation functions SN-203
 card reader SN-204
 console SN-203
 printer SN-204
 punch SN-205
 writing decompaction
 routines SN-223
 \$RJE2780 utility

commands CT-339
 description CT-337, OR-416
 host subsystems CT-338
 \$RJE3780 utility
 commands CT-339
 description CT-337, OR-418
 function SG-64
 host subsystems CT-338
 \$RMU
 See Remote Management Utility
 \$RMUPA CT-297
 \$\$ operator command
 entering commands and parameters CA-412
 function SG-61
 subcommands
 ALT CA-415
 DALL CA-418
 DE CA-419
 DG CA-420
 DISP CA-421
 HOLD CA-426
 KEEP CA-427
 REL CA-429
 STOP CA-430
 WRES CA-431
 WSTP CA-433
 WSTR CA-435
 \$SMJOB program OR-73
 \$SMLOG program OR-73
 \$SMMAIN program OR-46
 \$SMMPRIM, primary option menu OR-55
 \$SMM02, secondary option menu OR-57
 \$SMM03, secondary option menu OR-60
 \$SMM04, secondary option menu OR-62
 \$SMM05, secondary option menu OR-63
 \$SMM06, parameter selection menu OR-64
 \$SMM07, parameter selection menu OR-65
 \$SMM08, communication utilities option OR-67
 \$SMM09, secondary option menu OR-68
 \$SNADEF, SNA configuration file SN-103
 \$SPLUT1 utility
 commands
 CA CA-440
 CD CA-441
 DS CA-441
 EN CA-440
 GS CA-440
 MA CA-441
 MJ CA-441
 RS CA-443
 SO CA-441
 description OR-427
 functions CA-437, SG-66
 invoking CA-440
 \$SPOOL spool manager
 functions CA-400
 internals ID-135
 starting spool facility CA-401
 terminating spooling CA-402
 \$\$\$INIT utility
 description SG-441
 overview SG-436
 \$STOREMAP storage map ID-26
 \$\$SYSCOM system common data
 area ID-13, SG-172
 \$\$SYSLOG system logging
 device SG-116
 \$\$SYSLOGA alternate system logging
 device SG-116
 \$\$SYSLOGB alternate system logging
 device SG-116
 \$\$SYSRTR system printer SG-116
 \$SIASM Series/1 macro assembler
 data sets required PP-28
 description PP-27
 function SG-67
 internals ID-279
 output PP-38
 overview SG-7
 \$SISUT1 utility
 commands CT-101
 description CT-101, OR-425
 function SG-66
 invoking CT-101
 \$T operator command OR-36, SG-61
 \$TAPEUT1 utility
 commands OR-432
 description OR-431
 function SG-58, SG-63
 used in Version 3
 conversion SG-439
 \$TERMUT1 utility OR-462, SG-66
 \$TERMUT2 utility
 function SG-66
 process 4978 image or control store OR-468
 restore 4974 image OR-468
 \$TERMUT3 utility OR-487, SG-66
 \$TRAP utility OR-491, SG-65
 functions SG-65
 \$U operator command OR-37, SG-61
 \$UNPACK subroutine
 description CT-58
 example CT-59
 \$UPDATE utility
 completion codes MC-216, PP-115
 data sets required PP-105
 description OR-494, PP-105
 function SG-67
 output PP-115
 use considerations PP-116
 \$UPDATEH utility (host) PP-117, SG-67
 \$VARYOFF operator command
 description OR-38
 function SG-61
 syntax OR-38
 \$VARYON operator command
 description OR-40
 function SG-61
 syntax OR-40
 with standard labeled tape SG-296
 \$VERIFY utility
 example CA-127
 functions CA-124, OR-498, SG-65
 invoking CA-125, OR-499
 messages CA-134
 storage requirements CA-139
 \$W operator command OR-44, SG-62
 /*CONCAT control record SN-206
 /*END control record SN-206
 #1 index register 1 LR-7
 #2 index register 2 LR-7

A

A-conversion LR-169
A/I
 See analog input
A/O
 See analog output
AB command (\$S1S1UT1) CT-102
ABORT command (\$RJESNA) SN-210
ABORT command (\$RJE2780/\$RJE3780) CT-339
aborting Series/1 to Series/1 write CT-102
aborting task level (SVCABEND) ID-40
absolute record copying OR-101
accepting host initiated transactions SN-40
access level, BSC CT-214
Access Method, Indexed
 See Indexed Access Method
accessing system control blocks ID-331
acknowledgment, message SN-33
ACQUIRE parameter, NETINIT SN-21
action codes
 I/O error exit SM-143
 input exit SM-140
 output exit SM-142
ACTION, Multiple Terminal Manager CALL
 coding description CA-225
 internals M-10
 overview of CA-207
activating a stopped task OR-152
activating error logging, \$LOG OR-400
activating paging IG-167
activating TRAP function of storage dump OR-491
activation, session SN-17
active, task supervisor state ID-30
ACTLU command SN-134
ACTPU command SN-134
AD command (\$DICOMP) CT-166
AD command (\$PDS) SG-330
AD command (\$TERMUT2) OR-472
AD subcommand (\$DICOMP) CT-171
ADAPTER statement SG-90
ADD data manipulation instruction
 coding description LR-38
 overview LR-14
 precision table LR-39
adding a new device, Multiple Terminal Manager M-6
adding a null data set on tape volume OR-461
adding member to data base (graphics) CT-166
adding session manager options OR-80
adding support for new terminals
 coding GENxxxx ID-116
 coding the support module ID-114
 implementation ID-113
 modifying the TERMINAL statement ID-116
ADDRESS parameter, SNAPU SN-105
address relocation translator SG-43

addressing indexing feature LR-7
addroot sort
 See SORTA (record address sort)
ADDV data manipulation instruction
 coding description LR-38
 index register use LR-42
 overview LR-14
 precision table LR-42
AF command (\$INITDSK) OR-354
AI
 See analog input
AL command (\$DICOMP) CT-167
AL command (\$DISKUT1) OR-172
AL command (\$DIUTIL) CT-195
AL command (\$IAMUT1) IG-151
AL command (\$JOBUTIL) OR-378
ALIGN statement
 coding description LR-45
 overview LR-14
aligning spool job forms CA-415
ALLOCATE function, Remote Management Utility CT-258, ID-200
allocating a directory IG-69
allocating data set (\$RMU) CT-258
allocating fixed head volume OR-354
allocating journal data set SN-212
allocating member of partitioned data set SG-334 with \$DIUTIL command CT-195
allocating secondary index directory IG-151
allocating volume OR-353
 allocating OR-354
altering alternate tracks OR-136
altering member (graphics) CT-167
altering spool job printing CA-415
altering terminal configuration OR-462
alternate collating sequence statement
 See ALTSEQ (alternate collating sequence) statement
alternate sectors SG-49
ALTSEQ (alternate collating sequence) statement
 affect on control field contents SM-87
 ASCII and EBCDIC differences SM-88
 definition SM-8
 EBCDIC character codes SM-88
 example of use SM-8
 example, specifying SM-88
 how to specify SM-87
ALTSEQ statement, position descriptions
 position 1 - statement type SM-87
 position 3-4 - From character SM-88
 positions 35-80 - comments SM-90
 positions 5-6 - To character SM-88
 positions 7-34 - additional characters SM-90
analog input
AT \$IOTEST command OR-373
control block ID-162
description SG-357
IODEF statement SG-361

overview SG-58
 SBIO instruction SG-367
 analog output
 AO \$IOTEST command OR-368
 control block ID-162
 description SG-357
 IODEF statement SG-361
 SBIO instruction SG-367
 SENSORIO statement SG-101
 AND instruction
 coding description LR-46
 overview LR-14
 AO
 See analog output
 application program
 automatic initialization and
 restart SG-174
 introduction LR-1
 preparation of PP-1
 size estimating SG-225
 structure LR-2
 application program major
 node SN-323
 application program stub, Multiple
 Terminal Manager M-9
 application, identifying
 host SN-17
 arithmetic operators LR-4
 ascending/descending sequence
 option SM-60
 ASCII (American National Standard
 Code), changing SM-8
 ASCII terminal
 codes SG-114
 configuring SG-113
 definition CT-6
 devices supported SG-17
 graphics CT-151, SG-81
 ASMERROR instruction,
 \$EDXASM ID-256
 assembler
 See \$EDXASM and \$SIASM
 assigning \$RJESNA printer SN-213
 assigning alternate for defective
 4963 sector OR-136
 assigning DEFINE key in 4978
 control store OR-472
 asynchronous communications
 See terminal support
 AT command (\$DEBUG) OR-149
 ATTACH instruction
 coding description LR-48
 internals ID-35
 overview LR-28, SG-30
 attention event sequence,
 example SN-40
 attention handling, terminal
 I/O SG-62
 attention handling, terminal
 support
 internals ID-102
 See also interrupt servicing
 attention list
 processing for terminal
 support CT-20, CT-21
 See also ATTNLIST statement
 ATTNEV parameter, NETINIT SN-20
 ATTNLIST statement
 coding description LR-50
 overview LR-28, SG-30
 use in terminal support CT-21
 with a two task program LR-405
 ATTR define attribute characters,
 \$IMAGE command OR-326

attribute characters, 3101 CT-36,
 CT-38
 AUTOCALL control statement (\$EDX-
 LINK) PP-51
 AUTOCALL data set
 (\$EDXLINK) PP-74
 AUTOCALL data set record format
 (\$EDXLINK) PP-74
 autocall feature
 description PP-87
 statement requirement
 (WXTRN) LR-399
 AUTOCALL option (\$EDXLINK) PP-74
 AUTOCALL processing
 (\$EDXLINK) PP-75
 automatic initialization and
 restart SG-174
 automatic update indicator IG-69
 AV command (\$INITDSK) OR-353

B

B line command (\$FSEDIT) OR-280
 Backing up disk or diskette volume
 on tape OR-457
 backup IG-94
 backup dump restore utility,
 \$MOVEVOL OR-402
 base records, indexed file
 definition of CA-33, IG-61
 loading CA-33, IG-61
 base SNA function codes SN-21,
 SN-87
 basic exchange diskette
 data set copy utility OR-107
 exchanging between
 systems SG-54
 basic supervisor and emulator
 See supervisor/emulator
 batch job processing
 See \$JOBUTIL job stream proces-
 sor
 BCB (buffer control block inter-
 nals T-35
 BEEP subroutine CA-466
 BEEP, Multiple Terminal Manager
 CALL
 coding description CA-229
 internals M-9
 overview of CA-207
 BF command (\$IAMUT1) IG-144
 BID command SN-134
 binary synchronous communications
 access level CT-214
 automatic retry SG-20
 BSCAM/BSCAMU module
 descriptions ID-64
 BSCLINE statement SG-92
 conversational
 operations CT-216
 device data block (BSCDDB)
 description ID-167
 listing ID-340
 features of BSCAM CT-214,
 SG-19
 instruction formats ID-174
 instruction overview CT-219
 instruction processing
 (BSCAM) ID-177
 internals ID-167

multipoint operations CT-215, SG-19
 overview CT-10, SG-82
 point-to-point lines SG-19
 Remote Management Utility requirements CT-250
 return codes CT-246, MC-227
 sample programs CT-300
 special labels for ID-179
 storage requirements SG-220
 test utility, \$BSCUT2 CT-239
 trace printing utility, \$BSCUT1 CT-237
 trace utility, \$BSCTRCE CT-235
 transmission protocol ID-184
 utilities overview CT-213
 BIND checking SN-120
 BIND command SN-135
 BIND parameter control block, BND SI-282
 BIND parameters SN-121
 BIND post codes MC-218, SN-94
 BIND processing SN-142
 blank screen, \$B operator command OR-26, SG-61
 BLCC, block lock line count AM-10, IM-11
 BLDTXT subroutine, \$EDXASM ID-263
 BLINK subroutine CA-467
 BLINK, Multiple Terminal Manager CALL
 coding description CA-231
 internals M-10
 overview CA-207
 block and record locks AM-10, IM-11
 block header, Indexed Access Method description AM-6, IM-6
 equates AM-74
 blocked sequential IG-164
 BLP
 See also bypass label processing
 tape management utility OR-431
 BND, BIND parameter control block SI-282
 BOT (beginning-of-tape) LR-26
 bottom editor subcommand (\$EDIT1/N) OR-213
 boundary requirement, full-word DO LR-21
 IF LR-21
 PROGRAM LR-266
 BP command (\$DEBUG) OR-151
 BPIACB, block place header AM-10, IM-11
 bracket initiation, PLU SN-128
 bracket initiation, SLU SN-128
 bracket protocol SN-128
 brackets
 definition SN-127
 error conditions SN-130
 indicator rules SN-128
 initiation rules SN-128
 termination rules SN-129
 branch vector table, BVT SI-286
 BROWSE option (\$FSEEDIT) OR-265
 BSC
 See binary synchronous communications
 BSC 8-line control data rates CT-12
 use with #2094 CT-12

BSCAM
 See binary synchronous communications
 BSCCLOSE instruction coding description CT-219
 expanded format ID-174
 processing ID-178
 BSCDDB (binary synchronous device data block) description ID-167
 listing ID-340
 BSCEQU (BSCLINE equate table) LR-431
 BSCIA immediate action routine ID-178
 BSCIOCB statement CT-218, ID-174
 BSCLINE statement SG-92
 BSCOPEN instruction coding description CT-220
 expanded format ID-175
 processing ID-178
 BSCREAD instruction coding description CT-221
 example CT-234
 expanded format ID-175
 processing ID-178
 required buffers CT-220
 types of BSC read operations CT-223
 BSCWRITE instruction coding description CT-226
 example CT-233
 expanded format ID-176
 processing ID-178
 required buffers CT-220
 types of BSC write operations CT-227
 BSF (backward space file) LR-65, OR-449
 BSR (backward space record) LR-65, OR-449
 BTAM/BTAM-ES, channel attach considerations CT-418
 BTE (buffer table entry) description AM-6, IM-6
 equates AM-72, IM-114
 BU command (\$DIUTIL) CT-195
 BUFF parameter NETCTL SN-48
 NETGET SN-37
 NETPUT SN-30
 buffer
 allocation in Remote Management Utility CT-255
 areas, Multiple Terminal Manager M-23, M-37
 input, Multiple Terminal Manager CA-206
 management, terminal I/O ID-103
 output, Multiple Terminal Manager CA-206
 buffer control block (BCB) internals T-35
 BUFFER statement coding description LR-54
 overview LR-13
 buffer table entry (BTE) description AM-6, IM-6
 equates AM-72, IM-114
 buffers, specifying receive SN-114
 buffers, specifying send SN-114
 BUFNO parameter, SNAPU SN-106

BUFSIZ parameter, SNAPU SN-106
 BUILD option (\$EDXASM) PP-14
 building data member
 (graphics) CT-195
 bus
 See General Purpose Interface
 Bus
 BVT, branch vector table SI-286
 bypass label processing (BLP)
 description SG-300
 example SG-311
 BYTES parameter, NETGET SN-37
 BYTES parameter, NETPUT SN-30

C

CA attention request
 (\$COPYUT1) OR-113
 CA attention request
 (\$EDXASM) PP-13
 CA attention request
 (\$EDXLIST) PP-25
 CA attention request
 (\$FSEDIT) OR-270
 CA command (\$CHANUT1) CT-409
 CACB1 and CACB2 channel attach
 configuration statements CT-413
 CACLOSE instruction
 coding description CT-400
 overview CT-400
 return codes CT-442
 CAD command (\$COPYUT1) OR-113
 CADIOCB (channel attach device I/O
 control block)
 data map T-41
 description T-35
 CAIOCB (channel attach I/O control
 block) statement
 coding description CT-399
 data map T-59
 example CT-399
 internals T-35
 overview CT-399
 CALL (\$COPYUT1) OR-118
 CALL instruction
 coding description LR-58
 overview LR-20, SG-41
 program LR-58
 subroutine LR-58
 CALLFORT instruction
 coding description LR-61
 overview LR-20
 CANCEL command SN-135
 canceling a listing OR-202
 canceling dump OR-195
 canceling list option OR-270
 canceling program OR-27
 cancelling messages SN-52
 CAOPEN instruction
 coding description CT-400
 example CT-400
 overview CT-400
 return codes CT-440
 CAP command (\$COPYUT1) OR-113
 CAPCB (channel attach port control
 block)
 data map T-55
 description T-35
 CAPRNT instruction
 coding description CT-401
 example CT-402

 overview CT-401
 return codes CT-440
 card reader function,
 \$RJESNA SN-204
 CAREAD instruction
 coding description CT-402
 example CT-403
 illustration CT-389
 overview CT-402
 return codes CT-440
 CASTART instruction
 coding description CT-404
 example CT-404
 overview CT-404
 return codes CT-440
 CASTOP instruction
 coding description CT-405
 example CT-405
 overview CT-405
 return codes CT-440
 CATRACE instruction
 coding description CT-406
 example CT-406
 overview CT-406
 return codes CT-440
 CAWRITE instruction
 coding description CT-407
 example CT-407
 overview CT-407
 return codes CT-440
 CB command (\$DISKUT2) OR-182
 CC line command (\$FSEDIT) OR-280
 CCB (control block)
 Multiple Terminal Manager M-13
 CCB (terminal control block)
 description ID-115
 GPIB terminal control
 block ID-119
 listing ID-341
 spool extension ID-148
 CCBEQU (terminal control block
 equate table) LR-432
 CD command (\$COPY) OR-103
 CD command (\$TAPEUT1) OR-433
 CDATA, Multiple Terminal Manager
 CALL
 coding description CA-233
 internals M-9
 overview of CA-207
 CDRRM equate listing, Remote
 Management Utility CT-330
 CG command (\$COPYUT1) OR-116
 CH command (\$BSCUT2) CT-244
 CH command (\$GPIBUT1) CT-135
 CH command (\$UPDATEH) PP-120
 chain properties SN-126
 CHAIN supervisor routine ID-44
 chain, ECB/QCB/TCB ID-44
 CHAIND supervisor routine ID-44
 CHAINE supervisor routine ID-44
 CHAINP supervisor routine ID-44
 chains, definition SN-126
 chains, receiving SN-147
 chains, send operation SN-127
 chains, sending SN-144
 changing \$DEBUG base
 address OR-160
 changing character string
 with \$EDIT1/N OR-214
 with \$FSEDIT OR-273
 changing execution
 sequence OR-153
 changing GPIB partition CT-135

changing graphics or report
 display profile CT-165
 changing hardcopy device
 with \$BSCUT2 CT-244
 with \$TERMUT1 OR-467
 changing host library PP-120
 changing logical device
 name OR-466
 changing output volume PP-107
 changing partition
 assignment OR-29
 changing realtime data member
 name SG-332
 changing spool job
 forms code CA-415
 id heading CA-415
 number of copies CA-415
 output device CA-415
 changing tape label
 support OR-442
 changing terminal address assign-
 ment OR-465
 changing volume
 with CV (\$DISKUT1) OR-172
 with CV (\$DISKUT2) OR-181
 with CV \$UPDATE PP-107
 with CV \$UPDATEH PP-118
 with CV command
 (\$BSCUT1) CT-239
 changing 4978 key
 definition OR-473
 Channel Attach
 \$CAPGM CT-382
 BTAM considerations CT-418
 CAREAD illustration CT-389
 CAWRITE illustration CT-389
 closing a port CT-396
 control blocks
 data maps T-39
 descriptions T-35
 structure T-36
 device (4993) CT-380
 functions supported CT-379
 instruction expansions T-29
 instructions, listed CT-397
 issue I/O
 read CT-387
 write CT-388
 link module descriptions T-18
 link-editing CT-416
 module descriptions
 link modules T-18
 overlay programs T-32
 program modules T-8
 utility program
 modules T-26
 opening a port CT-386
 overlay program
 descriptions T-32
 overview CT-12, CT-379, SG-83
 post codes CT-435, MC-219
 preparing an application
 program
 assembling CT-416
 link-editing CT-416
 problem determination
 error logging CT-428
 online test (OLTEP) CT-434
 trace CT-419
 program module
 descriptions T-8
 return codes CT-437, MC-229
 sample programs CT-442
 starting a device CT-386
 stopping a device CT-396
 storage requirements SG-222
 system generation CT-412
 utility
 See \$CHANUT1 utility
 utility program module
 descriptions T-26
 character constants LR-85
 character image table OR-242
 \$FONT utility OR-242
 CHASE command SN-135
 CHGPAN, Multiple Terminal Manager
 CALL
 coding description CA-236
 internals M-9
 overview of CA-207
 CICS/VS considerations SN-172
 CICS/VS support SN-155
 CL primary command
 (\$FSEDIT) OR-275
 class interrupt vector
 table ID-11
 CLEAR command SN-135
 clear to send timeout SN-106
 clearing data set OR-181
 clearing screen OR-26
 CLEART subroutine CA-468
 CLOSET subroutine CA-469
 CLSRU close tape data set LR-65
 cluster controller SI-1
 clustered records IG-18
 CM command (\$COPYUT1) OR-115
 CM command (\$DIUTIL) CT-197
 CMDEQU LR-432
 CMDSETUP emulator entry
 point ID-14
 CNCTYPE parameter, SNAPU SN-105
 CNG command (\$COPYUT1) OR-116
 CO command
 (\$RJE2780/\$RJE3780) CT-341
 COBOL
 estimating storage
 requirements SG-228
 execution requirements SG-24
 overview SG-9
 program preparation require-
 ments SG-24
 use with Multiple Terminal
 Manager CA-283
 code translation
 tables ID-114
 terminal I/O ID-103
 code words, task LR-9
 coding the TERMINAL
 statement SG-116
 collating sequence
 changing SM-8
 specifying in header SM-67
 COLS line command
 (\$FSEDIT) OR-282
 COMMAND command (\$RJESNA) SN-211
 COMMAND command
 (\$RJE2780/\$RJE3780) CT-340
 command table, emulator
 description ID-14
 emulator setup routine ID-52
 listing ID-323
 commands processed by Series/1
 SNA SN-132
 comment control statement (\$EDX-
 LINK) PP-51
 comments
 in comment statement SM-68

in other specification state-
 ments SM-76
 commercial application
 facilities CA-2
 commercial products SG-4
 common AM-23
 common area, \$EDXASM ID-236
 common area, Multiple Terminal
 Manager M-16, M-34
 common data area
 See \$SYSCOM system common data
 area
 common emulator setup routine
 command table ID-14
 description ID-52
 communication products SG-4
 communications error
 function ID-209
 communications support SG-19
 communications utilities
 \$BSCTRCE CT-235
 \$BSCUT1 CT-237
 \$BSCUT2 CT-239
 \$HCFUT1 CT-367
 \$PRT2780 CT-343
 \$PRT3780 CT-343
 \$RJE2780 CT-337
 \$RJE3780 CT-337
 \$RMU CT-249
 session manager OR-66
 summarized SG-63
 communications vector table
 description ID-12
 listings ID-319
 compacted data, processing SN-223
 compaction SN-222
 compiler
 See \$EDXASM Event Driven
 Language compiler
 compiler common area,
 \$EDXASM ID-236
 completion codes
 \$EDXASM MC-208, PP-21
 \$EDXLINK MC-209, PP-73
 \$IAMUT1 CA-121, IG-186, MC-210
 \$JOBUTIL MC-212
 \$LINK MC-213, PP-97
 \$UPDATE MC-216, PP-115
 See also post codes, return
 codes
 component definition SN-160
 component flow, Multiple Terminal
 Manager M-3
 compressed byte string CT-60
 compressing library OR-95
 compression SN-222
 CONCAT instruction
 coding description CT-153
 overview CT-153
 concatenating data sets SN-206
 concatenation loops,
 avoiding SN-215
 concurrent execution LR-28
 conditional force
 continuation SM-115
 data types allowed SM-109
 definition SM-118
 how to specify SM-118
 keep/drop from output
 record SM-64
 length, maximum
 in a field statement SM-109
 specifying in header SM-60
 total in include set SM-118
 number, maximum SM-118
 required entries in field
 statement SM-109
 rules for SM-115
 configuration statements SG-89
 configuration, system
 See system configuration
 configuring terminal OR-463
 connecting Series/1 support SI-3
 connecting to indexed file CA-54,
 IG-87
 consecutive frames,
 specifying SN-106
 console function, \$RJESNA SN-203
 control block and parameter tables
 list LR-431
 control block linkages, Indexed
 Access Method AM-9, IM-9
 control blocks
 accessing ID-331
 analog input ID-162
 analog output ID-162
 Binary Synchronous Communi-
 cations (BSCDDB) ID-340
 chaining ID-334
 Channel Attach T-37
 data set control block (DSCB)
 data maps T-39
 descriptions ID-88, M-13,
 M-17, T-35
 equates M-32
 listing ID-360
 device data block (DDB)
 description ID-84
 listing ID-357
 digital input/output control
 block ID-162
 EXIO device data block ID-153
 GPIB terminal control
 block ID-119
 including ID-334
 Indexed Access Method
 description IM-5
 equates IM-81
 Multiple Terminal Manager
 description M-13
 equates M-33
 printer spooling
 descriptions ID-143
 listings ID-371
 process interrupt control
 block ID-162
 Remote Management Utility
 description ID-191
 listing ID-366
 sensor I/O DDB listing ID-358
 sensor-based I/O control block
 (\$BIOCB) ID-159
 tape device data block (TDB)
 description ID-92
 listing ID-377
 task control block (TCB)
 description ID-32
 listing ID-375
 terminal control block (CCB)
 description ID-115
 listing ID-341
 control field
 data types allowed SM-110
 examples SM-14
 how to specify SM-126
 keep/drop from output record
 default drop with variable
 records SM-65

specifying in header SM-64
 when to drop SM-65
 length, maximum SM-110
 in a field statement SM-60
 specifying in header SM-117
 total in include set SM-117
 number, maximum SM-105
 purpose SM-105
 required entries in field
 statement SM-109
 rules SM-115
 types SM-105
 control flow
 Indexed Access Method
 description AM-5
 equates AM-51
 with data paging IM-33
 control flow, access method AM-1,
 IM-1
 CONTROL IDCB command LR-201,
 SG-58
 control keys for test
 editors OR-202
 control operation internals,
 GPIB ID-121
 control operations, NETCTL SN-49
 CONTROL option (\$EDXASM) PP-14
 control records, \$LINK PP-88
 CONTROL tape instruction LR-64
 control, device instruction
 level LR-18
 controlling echo mode CA-109,
 IG-160
 controlling message
 exchange SN-9, SN-47
 conversational operations,
 BSC CT-216
 conversion algorithm for
 graphics ID-225
 conversion module (EBFLCVT) ID-65
 conversion, terminal I/O
 binary/EBCDIC ID-104
 converting floating
 point/binary ID-227
 converting program modules PP-7
 converting to Version 3
 \$MIGRID SG-441
 \$MIGRATE SG-457
 data SG-436
 CONVTB instruction
 coding description LR-70
 internal format ID-229
 overview LR-13
 return codes LR-70, MC-232
 CONVTD instruction
 coding description LR-75
 internal format ID-229
 overview LR-13
 return codes LR-75, MC-232
 copy code library
 (\$EDXASM) ID-244
 COPY control statement
 (\$EDXLINK) PP-52
 COPY instruction
 coding description LR-80
 overview LR-20
 copying block of text OR-280
 copying data set OR-103
 copying data sets with
 allocation OR-111
 copying line of text OR-280
 copying member
 with \$COPYUT1 OR-112
 with \$DIUTIL CT-197
 copying text OR-216
 copying volume OR-105
 correlation table entry,
 CTE SI-290
 correlation table entry,
 format SN-124
 correlation tables SN-124
 count record, Remote Management
 Utility CT-253, CT-298
 CP command (\$DIUTIL) CT-196
 CP command (\$GPIBUT1) CT-135
 CR command (\$IAMUT1) CA-105
 creating character image
 tables OR-242
 creating indexed file CA-105
 creating secondary index IG-76
 creating source data set OR-266
 creating stand-alone dump
 diskette OR-120
 creating supervisor for another
 Series/1 SG-214
 creating task error exit
 routines SG-256
 cross-partition services
 DEQ LR-92
 description and
 examples LR-432
 ENQ LR-120
 MOVE LR-225
 overview LR-11
 POST LR-240
 READ LR-290
 WAIT LR-385
 WHEREAS LR-388
 WRITE LR-391
 CSECT list, supervisor SG-465
 CSECT statement
 coding description LR-82
 overview LR-20
 CT command (\$TAPEUT1) OR-435
 CT command (\$TERMUT1) OR-463
 CTE parameter, SNALU SN-115
 CTE, correlation table
 entry SI-290
 customized system
 generation SG-194
 CV command (\$BSCUT1) CT-237
 CV command (\$COPY) OR-105
 CV command (\$DISKUT1) OR-172
 CV command (\$DISKUT2) OR-181
 CV command (\$UPDATE) PP-107
 CV command (\$UPDATEH) PP-118
 cycle steal status word (CSSW)
 description SG-267
 finding SG-268
 CYCLE, Multiple Terminal Manager
 CALL
 coding description CA-238
 internals M-10
 overview CA-206
 cylinder SG-46
 cylinder track sector (CTS)
 disk SG-49
 diskette SG-52

D

D/I See digital input
D/O See digital output
DACTLU command SN-136
DACTPU command SN-136
DAT, SNA data area SI-294
data
 definition statements LR-13
 floating-point LR-15
 floating-point arithmetic instructions
 data representation LR-16
 operations using index registers LR-16
 return code introduction LR-16
 formatting functions
 instructions LR-13
 module names LR-14
 integer LR-14
 integer and logical instructions LR-14
 manipulation instructions
 data representation LR-15
 overflow LR-15
 vector LR-15
 overview LR-13
 terminal input LR-32
 terminal output LR-31
Data Base/Data Communication (DB/DC) SI-1
Data Collection Interactive (DCI)
 overview SG-13
 storage requirements SG-222
data communication
 architecture SI-1
data control block, DCB SI-302
data exchange SG-23
data fields
 definition SM-105
 determining output record format and content SM-105
 how to specify SM-124
 rules SM-125
data files for \$SIASM ID-280
data formatting return codes MC-232
data interchange SG-23
data link buffer, DLB SI-317
data link terminal, DLT SI-327
data management SG-45
data management utilities
 \$COMPRES OR-95
 \$COPY OR-100
 \$COPYUT1 OR-111
 \$DASDI OR-119
 \$DISKUT1 OR-168
 \$DISKUT2 OR-180
 \$DISKUT3 SG-339
 \$INITDSK OR-343
 \$MOVEVOL OR-404
 \$PDS SG-321
 \$TAPEUT1 OR-431
 session manager OR-59
 summarized SG-63
data manipulation instructions
 ADD LR-38
 ADDV LR-41
 coding description LR-38
 DIVIDE LR-99
 MULTIPLY LR-230
 overview LR-14
 SUBTRACT LR-332
data manipulation, vector LR-15
data paging control block IM-6
data paging, Indexed Access Method
 activating IG-167
 control flow IM-33
 deactivating IG-166
 description IG-206
 extracting statistics IG-108
 getting paging statistics IG-169
 module descriptions IM-53
 page preemption IM-40
 problem determination IM-78
 selecting IG-167
 setting paging area size IG-168
 using IG-208
data record contents, text editor ID-387
data record, Remote Management Utility CT-253, CT-298
data representation, terminal I/O LR-31
data set allocation/deletion
 with \$DISKUT1 OR-172
 with \$DISKUT3 SG-339
 with \$JOBUTIL OR-378
 with \$PDS SG-322
 with session manager OR-47
data set concatenation SN-206
data set control block (DSCB)
 description ID-88
 listing ID-360
data set definition statement
 See DSD statement
data set description, DSD SI-334
data set directory
 internals ID-77
data set format
 \$FSEDIT OR-260
 \$PDS SG-323
 \$PRT2780 CT-343
 \$PRT3780 CT-343
data set naming
 conventions SG-283
data set ready (DSR) ID-178
data set ready timeout SN-105
data set utilities
 See data management utilities
data sets
 defining in DSD statement SM-71
 end-of-data requirements SM-82
 input SM-9
 output SM-9
 purposes SM-9
 required SM-12
 specification SM-9
 work SM-9
DATA statement coding
 description LR-84
data terminal ready
 timeout SN-105
data transfer initiation, terminal I/O ID-106
data transfer ready (DTR) ID-178
data types
 format, internal
 binary, signed SM-226
 character SM-222

decimal packed SM-226
 decimal, sign attached leading SM-226
 decimal, sign detached leading SM-226
 decimal, sign detached trailing SM-226
 decimal, unpacked SM-226
 digit SM-223
 zoned SM-223
 in field statement SM-110
 in include/omit statement SM-97
 data-set-shut-down, Indexed Access Method CA-150, IG-215
 data, invalid in output SM-82
 date/time
 description SG-207
 display, \$W operator command OR-44
 set, \$T operator command OR-36
 DC statement
 coding description LR-84
 overview LR-13
 DCB statement
 coding description LR-87
 function SG-57
 overview LR-18
 DCB, data control block SI-302
 DCBNO parameter, SNAPU SN-106
 DCE (directory control entry)
 description ID-80
 listing ID-347
 DCI (Data Collection Interactive) SG-13
 DD command (\$GPIBUT1) CT-136
 DD command (\$S1S1UT1) CT-102
 DD line command OR-283
 DDB (disk device data block)
 description ID-84
 listing ID-357
 DDB, device descriptor block SI-306
 DDBEQU (device data block equate table) LR-431
 DDE (device descriptor entry) ID-86
 DE command (\$DISKUT1) OR-172
 DE command (\$DIUTIL) CT-198
 DE command (\$IAMUT1) IG-153
 DE command (\$JOBUTIL) OR-379
 deactivating paging IG-166
 deactivating the network (SNADACT) SN-100
 debugging \$EDXASM overlay programs ID-274
 debugging aids (see also diagnostic aids) SG-21
 debugging utility, \$DEBUG
 data necessary for debugging OR-144
 description OR-142
 features OR-143
 tips and techniques OR-163
 decompaction routines SN-222, SN-224
 DEFINEQ statement
 coding description LR-90
 overview LR-24
 defining a multifunction attachment SG-90
 defining an interrupt key on 4978 OR-481
 defining device name CT-102
 defining disk volumes SG-286
 defining diskette volumes SG-285
 defining end character (GPIB) CT-135
 defining GPIB device CT-136
 defining horizontal tabs OR-331
 defining image dimensions OR-328
 defining indexed file CA-106, IG-146
 defining null characters OR-332
 defining processor SG-103
 defining protected fields (\$IMAGE) OR-342
 defining tape volumes SG-304
 defining vertical tabs OR-339
 defining 3107 attribute characters OR-326
 definition statements, data LR-13
 DELETE function, Remote Management Utility CT-260, ID-200
 DELETE, Indexed Access Method CALL
 description CA-67, IG-101
 return codes CA-68, IG-102
 deleting all spool jobs CA-418
 deleting data set
 on tape OR-461
 with \$JOBUTIL OR-379
 with \$PDS SG-335
 with Remote Management Utility CT-262
 deleting directory entry IG-153
 deleting elements CT-167
 deleting generic spool jobs CA-420
 deleting indexed file records CA-59, IG-93
 deleting member
 with \$DISKUT1 OR-172
 with \$DIUTIL CT-197
 deleting spool job CA-419
 deleting text
 with \$EDIT1/N OR-218
 with line (\$FSEDIT) OR-283
 deleting volumes, DV \$INITDSK command OR-365
 delimited data field
 data types allowed SM-109
 examples SM-18
 how to specify SM-124
 length, minimum/maximum SM-110
 required entries in field statement SM-109
 rules for specifying SM-115
 when to use SM-60
 density for diskette initialization, \$DASDI utility OR-122
 DEQ instruction
 coding description LR-92
 internals ID-49
 overview LR-28, SG-43
 supervisor function ID-36
 DEQBSC BSC routine ID-178
 DEQT instruction
 coding description LR-94
 function SG-79
 internals ID-110
 overview LR-30
 printer spooling ID-140
 descending sequence
 affect on control field contents SM-65
 option SM-60
 DETACH instruction
 coding description LR-97

internals ID-35
overview LR-28, SG-30
detached, task supervisor
state ID-30
device busy, resetting LR-136
device control block, spool (SPW)
description ID-149
listing ID-371
device data block
disk (DDB)
description ID-84
listing ID-357
sensor I/O listing ID-358
tape (TDB)
description ID-92
listing ID-377
device descriptor block,
DDB SI-306
device descriptor entry
(DDE) ID-86
device independence for terminal
support
between display stations and
printers CT-7, CT-15, SG-78
between 4978/4979 and
3101 CT-7, CT-31, CT-38
coding EDL instructions
for CT-39
for static screens CT-38
using the \$IMAGE
subroutines CT-41
device interrupt handling,
EXIO ID-156
device test utility,
\$IOTEST OR-368
device vector table ID-12
DF command (\$IAMUT1) CA-106,
IG-146
DI
See digital input
DI command (\$IAMUT1) CA-108,
IG-149
diagnosing
I/O errors SG-267
IPL problem SG-270
diagnostic aids
access method AM-31, IM-57
description SG-231
hardware-oriented SG-21,
SG-241
software trace table SG-249
software-oriented SG-22,
SG-236
summary SG-22
diagnostic utilities
\$DEBUG OR-142
\$DUMP OR-193
\$IOTEST OR-368
\$LOG OR-400
\$TRAP OR-491
overview SG-65
session manager OR-68
digital input
direct output to another
device, \$PDS utility SG-329
direct output, \$DICOMP subcom-
mand CT-171
external sync, XI \$IOTEST
command OR-371
I/O control block ID-162
IODEF statement SG-361
IOTEST command OR-370
overview SG-58
SBIO instruction SG-367
SENSORIO statement SG-101
digital output
I/O control block ID-162
IODEF statement SG-361
overview SG-58
SBIO instruction SG-367
direct access device
organization SG-53
direct graphics output CT-171
direct I/O
Series/1 to Series/1 CT-85,
LR-367
with IOCB LR-209
with PRINTTEXT LR-247
directory allocation IG-69
directory control entry (DCE)
description ID-80
listing ID-347
directory entries IG-71, SG-323
directory free space entry
(FSE) ID-82
directory member entry
(DME) ID-81
directory of instruction
locations LR-34
DISCONN Indexed Access Method CALL
description CA-69, IG-104
return codes CA-69, IG-104
DISCONNECT Multiple Terminal
Manager utility CA-306
disconnecting from indexed
file CA-54, IG-87
disk data sets
end-of-data requirements SM-82
rules for specifying SM-74
disk I/O instructions LR-16
DISK statement
coding SG-95
description SG-95
function SG-57
disk volumes, defining SG-286
disk/diskette capacity SG-285
disk/diskette directory
equates ID-347
disk/diskette return
codes LR-295, MC-234
diskette volumes, defining SG-285
disks/diskettes
converting to Version 3 SG-436
device data block (DDB)
description ID-84
listing ID-357
fixed-head SG-16, SG-48
logical volume SG-47
physical volume SG-46
storage capacity SG-285
symbolic addressing LR-10
utilities
\$COMPRES OR-95, SG-63
\$COPY OR-100, SG-63
\$COPYUT1 SG-63
\$DASDI OR-119, SG-63
\$DISKUT1 OR-168, SG-63
\$DISKUT2 OR-180, SG-63
\$IAMUT1 SG-63
\$INITDSK OR-343, SG-63
\$MOVEVOL OR-404, SG-63
\$TAPEUT1 SG-63
volume SG-16
disks/diskettes DISP command
(\$FONT) OR-244
dispatch queue, access
method AM-7, IM-7
display control member SG-324

display control member
format SG-326
display operator commands SN-212
display processor composer,
\$DICOMP CT-165
display processor interpreter,
\$DIINTR CT-185
display processor utilities,
general description CT-151
display processor utility,
\$DIUTIL CT-192
display profile elements,
\$PDS SG-326
display profiles, defining (graph-
ics) CT-164
displaying character image
tables OR-244
displaying file definition paramet-
ers IG-161
displaying indexed file
parameters CA-108, IG-149
displaying report line
items SG-330
displaying screen image OR-334
displaying spool status CA-421
displaying status of all
tasks OR-162
displaying storage OR-30, SG-61
displaying storage or
registers OR-154
displaying time and date SG-62,
SG-332
displaying variable SG-328
displaying 4978 program function
keys OR-410
Distributed Presentation Manage-
ment
advantages of DPM SN-164
MFS DPM compatibility SN-164
screen off-loading SN-164
screen presentation SN-162
SLU type P without DPM SN-165
DIVIDE instruction
coding description LR-99
overview LR-14
precision table LR-99
DLB, data link buffer SI-317
DLT, data link terminal SI-327
DME (directory member entry)
description ID-81
listing ID-348
updated by SETEOD SG-354
DO
See digital output
DO command (\$IOTEST) OR-370
DO instruction
coding description LR-102
overview LR-21
double-precision
floating-point
arithmetic LR-16
integer and logical LR-14
DOWN subcommand (EDIT/N) OR-219
DP command (\$BSCUT1) CT-237
DP command (\$TAPEUT1) OR-437
DPACTBUF decompression
buffer SN-223
DPACTRJE decompression
routine SN-223
DR command (\$IAMUT1) IG-150
DR function (\$PDS) SG-328
DR subcommand (\$DICOMP) CT-172
drawing line relative SG-331
drawing line, \$DICOMP
subcommand CT-178
drawing symbol CT-172
DS command (\$JOBUTIL) OR-380
DSCB (data set control block)
statement
coding description LR-106
disk/diskette
description ID-88
listing ID-360
function SG-57
overview LR-16
tape internals ID-93
Version 3 conversion consider-
ations SG-435
DSCBEQU (data set control block
equate table) LR-431
DSD statement
input SM-72
output SM-72
purpose SM-71
work SM-72
DSD statement, position
descriptions
position 1 - statement
type SM-72
position 12-17 - volume
name SM-74
position 2 - data set
type SM-72
position 20-24 - record
size SM-75
position 25-29 - block
size SM-75
position 3 - device SM-73
position 35-80 -
comments SM-76
position 4-11 - data set
name SM-74
positions 18-19 - record
format SM-74
DSD, data set description SI-334
DSECT LR-431
DSOPEN subroutine SG-347
DSR (data set ready) ID-178
DTR (data transfer ready) ID-178
dump and restore volume OR-402
DUMP function, Remote Management
Utility CT-262
dumping storage partition
(\$RMU) CT-262
dumping to printer
with \$DICOMP CT-168
with \$DUMP OR-194
with \$TAPEUT1 OR-437
dumping to terminal
with \$DICOMP CT-168
with \$DISKUT2 OR-182
dumping trace file on printer
(\$BSCUT1) CT-237
dumping trace file on terminal
(\$BSCUT1) CT-237
dumping trace table SG-250
DV command (\$INITDSK) OR-365
dynamic storage, specifying (\$EDX-
LINK) PP-77, PP-81

E

- E-conversion (Ew.d) LR-167
- EBCDIC (extended binary coded decimal interchange character)
 - changing SM-9
 - collating sequence SM-224
 - default SM-67
 - specifying in header SM-66
- EBFLCVT, EBDIC to floating-point conversion
 - internals ID-227
 - module description ID-65
- EC command (\$IAMUT1) CA-109, IG-160
- EC command (\$S1S1UT1) CT-102
- ECB address SN-20
- ECB statement
 - coding description LR-108
 - internals ID-45
 - overview LR-28, SG-30
 - with SBIOCBs ID-161
- echo test, Series/1 to Series/1 CT-102
- EDIT command (\$EDIT1/N) OR-204
- EDIT command (\$FONT) OR-246
- EDIT command (IMAGE) OR-328
- edit data set subroutine examples, text editor ID-388
- EDIT option (\$FSEDIT) OR-266
- editor subcommands, \$EDIT1/N OR-212
- EDL (Event Driven Language)
 - estimating storage requirements SG-225
 - features summarized SG-20
 - instruction format ID-55
 - interpreter module description (EDXALU) ID-66
 - operation codes ID-56
 - purpose LR-1
 - use with Multiple Terminal Manager CA-281
- EDXFLOAT module description ID-66
- EDXINIT module
 - in IPL ID-5
 - module description ID-66
- EDXLIST host listing
 - formatter OR-297, PP-41
- EDXSTART module description ID-67
- EDXSVCX/EDXSVCXU module
 - description ID-67
- EDXTIO terminal I/O
 - internals ID-97
 - module description ID-68
- EF command (\$IAMUT1) IG-161
- EJECT command (\$JOBUTIL) OR-381
- EJECT statement
 - coding description LR-110
 - overview LR-19
- ejecting printer page OR-31
 - description OR-31
- ELSE instruction LR-111
- emulator
 - See supervisor/emulator
- emulator command table
 - description ID-14
 - listing ID-323
- emulator setup routine
 - command table ID-14
 - internal description ID-52
 - listing ID-323
- emulator/interpreter functional flow ID-52
- ENCODE parameter, SNAPU SN-105
- END command (\$EDXASM) PP-13
- END command (\$EDXLIST) PP-25
- END command (\$FSEDIT) OR-275
- END command (\$IMAGE) OR-329
- END control record (\$LINK) PP-92
- END control statement (\$EDXLINK) PP-53
- end display (graphics) CT-177
- end of job, DOJ command (\$JOBUTIL) OR-382
- end of nested procedure, EOP command (\$JOBUTIL) OR-382
- END statement
 - coding description LR-112
 - overview LR-28
- end-of-data processing SG-287
- end-of-file, indicating with SETEOD SG-354
- end-of-transaction, coding SN-32
- ENDATTN instruction
 - coding description LR-113
 - overview LR-28, SG-30
- ENDDO instruction
 - coding description LR-114
 - overview LR-21
- ENDIF instruction
 - coding description LR-115
 - overview LR-21
- ENDPROG statement
 - coding description LR-116
 - overview LR-28, SG-30
- ENDRJE command (\$RJESNA) SN-211
- ENDSEQ Indexed Access Method CALL
 - description CA-71, IG-106
 - return codes CA-71, IG-106
- ENDSPPOOL command (\$RJE2780/\$RJE3780) CT-340
- ENDTASK instruction
 - coding description LR-118
 - overview LR-28
- ENQ instruction
 - coding description LR-120
 - internals ID-48
 - overview LR-28, SG-43
 - supervisor function ID-36
- ENQT instruction
 - coding description LR-122
 - function SG-79
 - internals ID-109
 - overview LR-30
 - printer spooling ID-139
 - use for virtual terminals CT-205
 - use in GPIB CT-110
 - use with logical screens CT-33
 - using to reserve a terminal CT-17
- enqueue, task supervisor function
 - See ENQ instruction
- entering and editing source statements OR-204, SG-64, SG-70
- entering command parameters and operands
 - prompt-reply mode OR-22
 - single line format OR-23
- entering unprotected output fields using \$IMAGE OR-342
- entry points, supervisor SG-465
- ENTRY statement
 - coding description LR-125
 - overview LR-20

entry, correlation table SN-124
 EOF1 tape label SG-295
 EOJ command OR-382
 EOP command (\$JOBUTIL) OR-382
 EOR instruction
 coding description LR-127
 overview LR-14
 EOT (end-of-tape) LR-26
 EOT parameter, NETPUT SN-30
 EOVI tape label SG-295
 EP subcommand (\$DICOMP) CT-177
 EQ (equal) LR-21
 EQU statement
 coding description LR-129
 overview LR-13
 equate listings AM-51
 equate tables
 accessing ID-331
 BSCDDB (BSC device data block) ID-340
 CCB (terminal control block) ID-341
 communication vector table ID-346
 DDB sensor I/O device data block ID-358
 DDB, disk/diskette data block ID-357
 disk/diskette directory ID-347
 DSCB (data set control block) ID-360
 emulator command table ID-323 including ID-334
 Indexed Access Method IM-81
 Multiple Terminal Manager M-31 parameter and control block LR-431
 printer spooling ID-371
 program header listing ID-365
 referencing ID-23
 RCB (Remote Management Utility) ID-366
 Remote Management Utility general equates ID-369
 TCB (task control block) ID-375
 TDB (tape device data block) ID-377
 ERASE instruction
 coding description LR-131
 function SG-79
 overview LR-30
 ERRCODE parameter, NETINIT SN-20
 ERREXIT (load mode) IG-119
 ERREXIT (process mode) IG-123
 error codes
 See return codes
 error exit I/O SM-68
 error exit, task
 \$RJESNA SN-217
 Series/1 SNA SN-11
 error handling
 GPIB CT-131
 I/O error logging OR-400, SG-257
 Indexed Access Method error exit CA-148, IG-213
 PROGRAM (TERMERR=) LR-270
 program checks SG-273
 PSW (processor status word) SG-275
 Remote Management Utility CT-327
 software trace SG-249
 system-supplied SG-253
 TASK (TERMERR=) LR-336
 task error exit SG-253
 terminal I/O LR-30
 error handling, \$RJESNA SN-216
 error information, extended SN-87
 error logging (Indexed Access Method) IG-216
 error logging, channel attach CT-428
 error logging, SDLC SN-12
 error messages, Sort/Merge SM-247
 error, I/O, header option SM-66
 ERRORDEF LR-432
 ERRORS command (\$EDXLIST) PP-25
 errors option (\$EDXASM) PP-13
 errors, debugging Sort/Merge SM-179
 establishing a session, overview SN-6
 estimating storage
 See storage estimating
 event control block
 See ECB statement
 Event Driven Language
 See EDL (Event Driven Language)
 events, specifying attention SN-20
 EX command (\$TAPEUT1) OR-439
 example of Series/1 SLU type P/DPM configuration
 host components SN-170
 Series/1 control application SN-171
 Series/1 screen formatting/presentation support SN-170
 Series/1 SLU type P support SN-170
 exchange station ID SN-107
 exchange station ID, format SN-302
 exchanging data SG-23
 EXEC command (\$JOBUTIL) OR-383
 EXEC function, Remote Management Utility CT-264, ID-204
 executing program
 EXEC function, \$JOBUTIL OR-383
 EXEC function, Remote Management Utility CT-264
 overview SG-33
 PASSTHRU function, Remote Management Utility CT-281
 SHUTDOWN function, Remote Management Utility CT-276
 utilities (session manager) OR-65
 executing, task supervisor state ID-30
 exercising GPIB devices OR-292
 exercising tape OR-439
 exhibiting file definition parameters IG-161
 EXIO instruction
 coding description LR-136
 function SG-57
 overview LR-18, SG-20
 return codes LR-141, MC-236
 storage requirements SG-220
 EXIO support internals ID-153
 EXIODDB device data block ID-153
 EXIODEV considerations SN-103
 EXIODEV statement SG-98

exit routines
 examples SM-146
 I/O error SM-143
 input SM-137
 output SM-139
 when to use SM-33
 with vs. without synchroniza-
 tion SM-131
 writing SM-137
 EXOPEN instruction
 coding description LR-139
 function SG-57
 interrupt codes LR-141
 overview LR-18
 return codes LR-141
 extended error information SN-87
 extended error information,
 requesting SN-20
 extended error return codes
 BIND event MC-218, SN-94
 NETBIND MC-260, SN-94
 NETCLOSE MC-261, SN-93
 NETOPEN MC-262, SN-89
 NETRECV MC-263, SN-90
 NETSEND MC-264, SN-91
 NETUBND MC-267, SN-94
 READ MC-233, SN-95
 WRITE MC-233, SN-96
 extended summary field
 example SM-29
 how to specify SM-125
 length, maximum SM-110
 when to use SM-29
 external I/O level support inter-
 nals ID-153
 external sync DI/DO, XI/XO, \$IOT-
 EST command OR-368
 EXTRACT Indexed Access Method CALL
 description CA-73, IG-108
 return codes CA-74, IG-109
 EXTRN statement
 coding description LR-145
 overview LR-20

F

F-conversion (Fw.d) LR-166
 factor 1 field SM-99
 factor 2 constant SM-102
 factor 2 field SM-101
 FADD instruction
 coding description LR-147
 overview LR-14
 return codes LR-148, MC-238
 FAN Multiple Terminal Manager
 CALL CA-240
 fast path feature (IMS/V5) SN-161
 FAT, function address
 table SI-337
 FCA (file control area) CA-243
 FCB (file control block)
 accessing CA-159
 contents CA-159
 description IM-5
 equates IM-93
 queue description IM-8
 FCBEQU (Indexed Access Method copy
 code module) LR-432
 FCBEQU module CA-12, IG-4
 FDIVD instruction
 coding description LR-150
 overview LR-14
 return codes LR-151, MC-238
 FE command (\$HCFUT1) CT-370
 fetching record, \$PDS
 utility SG-335
 field statement
 controlling sort
 sequence SM-105
 determining output record
 content SM-105
 determining output record
 format SM-105
 examples SM-14
 how to specify
 control fields SM-116
 data fields SM-124
 position summary SM-107
 required entries in field
 statement SM-109
 rules SM-112
 rules for control
 fields SM-121
 rules for data fields SM-128
 field statement, position
 descriptions SM-106
 length, minimum/maximum
 in a field statement SM-112
 position 1 - statement
 type SM-106
 position 13 - forced data char-
 acter SM-114
 position 14 - force continua-
 tion indicator SM-115
 position 2 - field type SM-106
 position 3 - data type SM-109
 positions 12-13 -
 overflow/nonoverflow
 indicator SM-114
 positions 12-13 - test and
 insert characters SM-113
 positions 15-16 - extended
 summary field SM-115
 positions 35-80 -
 comments SM-115
 positions 4-11 - field location
 start and end SM-112
 rules for total in include
 set SM-117
 field table (FTAB)
 \$IMDATA subroutine CT-53
 \$IMPROT subroutine CT-51
 format of CT-51
 file
 backup IG-94
 backward space file
 (BSF) LR-65
 definition LR-26
 forward space file (FSF) LR-65
 sequential blocked IG-164
 sequential unblocked IG-164
 tape control commands LR-64
 file control block (FCB)
 accessing CA-159
 contents CA-159, IG-26
 description IM-5
 equates IM-93
 extracting IG-93, IG-108
 location CA-18, IG-40
 queue description IM-8
 file control block extension
 accessing CA-159
 contents CA-157, IG-40
 extracting IG-93, IG-108
 location CA-18, IG-40

file locking system IM-7
file manager, Multiple Terminal Manager M-8
file table, Multiple Terminal Manager M-19, M-35
FILEIO, Multiple Terminal Manager CALL
 coding description CA-242
 internals M-9
 overview of CA-208
 return codes CA-247
FIND command OR-276
FIND instruction
 coding description LR-153
 overview LR-21
FIND subcommand (\$EDIT1/N) OR-221
FINDNOT instruction
 coding description LR-156
 overview LR-21
finite state machine (FSM), description SN-127
FIRSTQ instruction
 coding description LR-159
 overview LR-24, SG-42
fixed blocked record format SM-79
fixed storage area ID-11
fixed unblocked record format SM-77
fixed-head devices SG-48
floating-point
 arithmetic instructions LR-15
 binary conversions ID-227
 EDXFLOAT/NOFLOAT module description ID-66
 return codes LR-16, MC-238
FM profiles 3 and 4
 differences SN-117
 supported commands SN-117
FMH parameter, NETPUT SN-30
FMULT instruction
 coding description LR-160
 overview LR-14
 return codes LR-161, MC-238
force-all
 See force-otherwise
force-otherwise
 data types allowed SM-120
 definition SM-120
 keep/drop from output record SM-65
 length, maximum
 in a field statement SM-112
 specifying in header SM-60
 total in include set SM-120
 number, maximum SM-120
 required entries in field statement SM-109
 rules for SM-121
forced control field
 conditional force
 continuation SM-115
 definition SM-119
 how to specify SM-119
 data types allowed SM-109
 default drop with variable records SM-65
 force-otherwise
 definition SM-117
 how to specify SM-117
 keep/drop from output record
 default drop with variable records SM-65
 example SM-14
 specifying in header SM-64
 when to drop SM-65
 length, maximum
 in a field statement SM-112
 specifying in header SM-60
 total in include set SM-117
multiconditional force
 definition SM-117
 how to specify SM-117
 number, maximum SM-117
 required entries in field statement SM-109
 rules for SM-121
 unconditional force
 definition SM-117
 example SM-22
 how to specify SM-117
forced data character
 examples SM-22
 forced control field, contrast SM-124
 how to specify SM-124
 required entries in field statement SM-109
 rules for specifying SM-128
format
 illustrated (program) LR-6
 instruction (general) LR-3
FORMAT statement
 A-conversion LR-169
 alphameric data LR-168
 blank lines in output LR-171
 coding description LR-163
 conversion of alphameric data LR-169
 conversion of numeric data LR-164
 E-conversion LR-167
 F-conversion LR-166
 formatted screen images OR-321
 function SG-79
 H-conversion LR-168
 I-conversion LR-165
 module names LR-14
 multiple field format LR-171
 numeric data LR-164
 overview LR-13
 repetitive
 specification LR-171
 storage considerations LR-172
 using multipliers LR-171
 X-type format LR-170
formatted screen return codes
 \$IMDATA subroutine MC-239
 \$IMOPEN subroutine MC-239
 \$IMPROT subroutine MC-240
formatting instructions,
 data LR-13
forms control
 for display screens CT-23
 for printers CT-24
FORTRAN IV
 estimating storage requirements SG-228
 execution requirements SG-26
 link editing SG-72
 overview SG-8
 program preparation requirements SG-25
 use with Multiple Terminal Manager CA-285
FPCONV instruction
 coding description LR-175
 overview LR-14
frame count, specifying SN-113

frames per transmission, considerations SN-111
 free space entry (FSE) ID-802
 FREEMAIN instruction ID-25
 FSE (directory free space entry) description ID-82
 listing ID-349
 FSF (forward space file) LR-65
 FSF (forward spacefile) OR-449
 FSR (forward space record) LR-65, OR-449
 FSUB instruction coding description LR-178
 index registers LR-179
 overview LR-14
 return codes LR-179, MC-238
 FTAB command (\$IMAGE) OR-330
 refid=pfkeys, listing, KEYS \$IMAGE command OR-331
 FTAB, Multiple Terminal Manager CALL coding description CA-254
 internals M-26
 overview of CA-207
 return codes CA-255
 full-word boundary requirement DO LR-21
 IF LR-21
 PROGRAM LR-266
 function address table, FAT SI-337
 function management headers, definition of SN-7
 function process subroutines adding new functions ID-214
 definition ID-189
 logic flow ID-196
 function table, Remote Management Utility ID-191

G

gather read operation CT-37, CT-40, CT-43
 GE (greater than or equal) LR-21
 general instruction format LR-3
 General Purpose Interface Bus \$GPIBUT1 utility commands CT-134
 description CT-134
 configuration CT-113
 cycle steal status CT-132
 data transfers CT-114
 device addresses CT-109
 device group operation CT-121
 device roles CT-108
 EDX support components CT-108
 error detection cycle steal status CT-133
 residual status block CT-133
 error handling CT-131
 function SG-66
 initialization CT-112
 internals control operations ID-121
 functional flow ID-120
 input operations ID-124
 output operations ID-125
 start I/O operation ID-125

terminal control block (CCB) ID-119
 interrupt handling CT-116
 interrupt status byte CT-132
 loading programs CT-110
 overview CT-7, CT-107, SG-18, SG-81
 parallel polling CT-118
 return codes MC-272
 roles of attached devices controller CT-108
 listener CT-108
 talker CT-108
 sample program CT-127
 serial polling CT-118
 service requests (SRQ) CT-116
 storage requirements SG-219
 system generation CT-131
 TERMCTRL coding description LR-371
 terminal I/O considerations CT-110
 translated data (XLATE=NO) CT-110
 universal unlisten CT-113
 user buffer CT-111
 generating a diskless system SG-214
 generating logical unit control block SN-114
 generating physical unit control block SN-104
 generating the supervisor SG-167
 GET command (\$FONT) OR-247
 GET Indexed Access Method CALL description CA-75, IG-111
 return codes CA-77, IG-113
 GETCUR subroutine CA-470
 GETCUR, Multiple Terminal Manager CALL coding description CA-258
 internals M-10
 overview CA-207
 GETDAT subroutine CA-471
 GETEDIT instruction coding description LR-181
 function SG-79
 overview LR-13
 return codes LR-183, MC-232
 GETFMT subroutine CA-473
 GETMAIN instruction ID-24
 GETPAR3 subroutine ID-53
 GETSEQ Indexed Access Method CALL description CA-78, IG-114
 return codes CA-80, IG-116
 GETTIME instruction coding description LR-188
 overview LR-32, SG-42
 time/date LR-188
 getting paging statistics IG-169
 GETVAL subroutine, \$EDXASM ID-261
 GETVALUE instruction coding description LR-190
 function SG-79
 overview LR-30
 GIN instruction coding description CT-154
 overview CT-154
 global area, \$EDXASM ID-246
 GLOBAL ATTNLIST LR-50
 GO command (\$DEBUG) OR-152
 GOTO command OR-153
 GOTO instruction coding description LR-196

- overview LR-21
- GP command (\$GPIBUT1) CT-136
- GPIB
 - See General Purpose Interface Bus
- GPRESUME command (\$GPIBUT1) CT-143
- graphics
 - conversion algorithm ID-225
 - functions CT-151
 - general description SG-81
 - hardware considerations CT-151
 - instructions
 - CONCAT CT-153
 - GIN CT-154
 - overview SG-81
 - PLOTGIN CT-155
 - XYPLOT CT-158
 - YTPLOT CT-159
 - overview CT-8
 - programming example CT-160
 - requirements CT-151
 - terminals CT-8, SG-81
 - utilities
 - \$DICOMP CT-165
 - \$DIINTR CT-186
 - \$DIUTIL CT-192
 - summarized OR-7, SG-65
- GT (greater than) LR-21

H

- H-conversion LR-168
- half-duplex flip-flop SN-131
- hardcopy function for terminals (PF6) ID-108, OR-9
- hardware features
 - for binary synchronous communications CT-11
 - for Remote Management Utility CT-252
 - for terminal support CT-5, CT-9
- hardware levels SG-30
- hardware options, SDLC SN-108
- hash table, data paging IM-34
- HASHVAL instruction
 - coding description LR-198
 - overview LR-14
- HCF
 - See Host Communications Facility
- HDR1 tape label SG-295
- header labels, tape SG-293
- header record, text editor ID-385
- header statement
 - examples SM-14
 - purpose SM-60
 - rules SM-60
- header, position descriptions
 - position 1 - statement type SM-60
 - position 11 - ascending/descending sequence SM-63
 - position 12 - message output SM-64

- position 13 - keep/drop control fields SM-64
- position 14 - warning message option SM-66
- position 15 - data set organization SM-66
- position 16 - collating sequence SM-67
- position 17 - I/O error processing SM-67
- position 35-80 - comments SM-68
- positions 2-7 - Sort/Merge application SM-60
- positions 8-10 - sum of control field lengths SM-62
- HELP command (\$DEBUG) OR-153
- hold line active timeout SN-106
- holding spool jobs(s) CA-426
- horizontal tabs, defining with \$IMAGE OR-331
- host assembler PP-39
- Host Communications Facility
 - data set
 - characteristics CT-348
 - host data sets CT-348
 - internals ID-181
 - overview CT-11, CT-347, SG-14, SG-83
 - programming examples CT-365
 - requirements CT-347
 - return codes CT-373, MC-276
 - storage and performance CT-352
 - system status data set CT-350
 - TP instructions
 - CLOSE CT-355
 - FETCH CT-356
 - OPENIN CT-357
 - OPENOUT CT-358
 - READ CT-359
 - RELEASE CT-360
 - SET CT-361
 - SUBMIT CT-362
 - TIMEDATE CT-363
 - WRITE CT-364
 - TPCOM module description ID-73
 - utility program (\$HCFUT1) CT-367
- host considerations, \$RJESNA
 - DOS/VSE (VSE/POWER) SN-235
 - OS/VS1 (RES) SN-233
 - OS/VS2 MVS (JES2) SN-227
 - OS/VS2 MVS (JES3) SN-231
 - VTAM NCP SN-225
- host data set naming conventions, Host Communications Facility CT-348
- host ID data list, building SN-17
- host initiated transactions, rejecting SN-41
- host logon, \$RJESNA SN-207
- host-initiated transactions, accepting SN-40
- HOSTCOMM statement SG-100
- HOSTID parameter, NETINIT SN-21
- HRJE command (\$RJESNA) SN-212
- HTAB command (\$IMAGE) OR-331
- HX subcommand (\$DICOMP) CT-177
- HXUT1, H exchange utility OR-298

I

I-conversion (Iw) LR-164
I/O control block spool (IOSPTBL)
 description ID-151
 listing ID-371
I/O device instruction
 level LR-18
I/O direct
 Series/1 to Series/1 CT-85,
 LR-367
 with IOCB LR-209
 with PRINTEXT LR-247
 with READTEXT LR-300
I/O error exit
 example SM-146
 in parameter list SM-133
 writing program with SM-137
I/O error logging
 activating SG-259
 data set list utility,
 \$DISKUT2 OR-180
 description SG-259
 device table ID-395
 equates ID-397
 invoking OR-400
 log control record ID-395
 log data set OR-400
LOG macro
 description SG-263
 syntax SG-263
 printing the errors SG-258
 record layout ID-397
 recording the errors ID-393,
 SG-258
 storage requirements SG-220
 tape log entries SG-262
USER instruction
 description SG-265
 example SG-266
 syntax SG-265
I/O error option SM-66, SM-67
I/O errors, SDLC device
 never retrieved SN-107
 retrieved until count
 exceeded SN-107
 retrieved until
 successful SN-107
I/O functions
 EXIO control internals ID-153
 HOSTCOMM statement SG-100
 instruction level
 control LR-18
 sensor-based
 sensor-based SBIOCB
 internals ID-159
 summarized SG-56, SG-59
 terminal SG-79
 timers SG-42
IACB (indexed access control
 block)
 creating CA-54, IG-87
 description IM-5
 equates IM-102
 queue description IM-8
IAM module CA-12
IAMEQU (Indexed Access Method copy
 code module) LR-432
IAMEQU module CA-12
IBM 3705 communication
 controller SI-1
ID command (\$INITDSK) OR-346
ID statement
 See \$ID statement
IDC, immediate device control
 block SI-340
IDCB statement
 coding description LR-201
 function SG-57
 IDCB command LR-201
 overview LR-18
IDCHECK function, Remote Manage-
 ment Utility CT-267, ID-203
identifying batch job stream
 continuation point OR-386
identifying batch job stream data
 set OR-380
identifying host
 application SN-17
IE command (\$IAMUT1) IG-155
IF instruction
 coding description LR-203
 overview LR-21
II command (\$INITDSK) OR-358
II line command (\$FSEDIT) OR-287
IIB (interrupt information
 byte) CA-220
IM function (\$PDS) SG-331
IM subcommand (\$DICOMP) CT-177
image dimensions, defining OR-328
image, formatted screen
 creating OR-321
 See also screen formats, termi-
 nal support
immediate action routines
 BSC access method
 (BSCIA) ID-178
 interrupt servicing ID-37
 specifying maximum
 number SG-104
 supervisor entry ID-38
immediate data LR-5
immediate device control block,
 IDC SI-340
immediate session
 termination SN-60
implied message acceptance SN-42
IMS/VS considerations SN-155
IMS/VS definition macros
 COMM macro SN-156
 NAME macro SN-160
 TERMINAL macro SN-157
 TYPE macro SN-157
IMS/VS support SN-155
IN command (\$DICOMP) CT-167
IN command (\$DIUTIL) CT-198
in-storage specifications SM-131
INCLUDE control statement (\$EDX-
 LINK) PP-53
INCLUDE control statement
 (\$LINK) PP-91
include set
 definition SM-93
 example SM-20
 how to specify SM-95
INCLUDE statement requirement
 (EXTRN) LR-145
include-all
 See include-remaining statement
include-remaining statement
 how to specify SM-95
 purpose SM-94
include/omit statement
 example SM-20
 how to specify SM-95
 purpose SM-93

include/omit, position
 descriptions
 factor 1 field, rules SM-99
 factor 2 constant,
 rules SM-102
 factor 2 field, rules SM-101
 position 1 - statement
 type SM-97
 position 14 - constant or
 field SM-101
 position 2 - AND/OR
 indicator SM-98
 position 3 - data types SM-98
 positions 12-13 - test rela-
 tional SM-101
 positions 15-22 - factor 2
 field SM-101
 positions 15-34 - factor 2
 constant SM-102
 positions 35-80 -
 comments SM-103
 positions 6-11 - factor 1 field
 location SM-99
 specifying in header SM-66
 independent indicator IG-68
 index block
 description CA-159
 internal description IM-6
 index entry, text editor ID-386
 indexed access control block
 description IM-5
 equates IM-102
 See also IACB (indexed access
 control block)
 Indexed Access Method
 \$IAM load module CA-189
 \$IAMUT1 utility CA-103, IG-141
 base records
 definition of CA-33, IG-61
 load mode IG-61
 loading CA-33, IG-61
 processing mode IG-61
 BTE (buffer table entry) IM-6
 buffer management CA-143,
 IG-210
 coding example CA-97, IG-233
 components CA-12, IG-4
 concepts and facilities CA-11,
 IG-1
 control block
 descriptions IM-5
 equate listings IM-81
 linkages IM-9
 control flow IM-1
 data block header IM-6
 data paging internals IM-33
 delete threshold
 parameter CA-118
 devices supported by CA-11,
 IG-3
 diagnostic aids IM-57
 dynamic file facilities CA-20,
 IG-41
 error recovery
 common IM-45
 data paging IM-53
 data-set-shut-down CA-150,
 IG-215
 deadlocks CA-151, IG-218
 error exit facility CA-148,
 IG-213
 long-lock-time CA-151,
 IG-218
 secondary index IM-55

system function return
 codes CA-149, IG-215
 features CA-2, IG-1
 file control block (FCB) IM-5
 file locking system IM-7
 I/O requests
 DELETE CA-67, IG-101
 DISCONN CA-69, IG-104
 ENDSEQ CA-71, IG-106
 EXTRACT CA-73, IG-108
 GET CA-75, IG-111
 GETSEQ CA-78, IG-114
 LOAD CA-81, IG-118
 overview CA-13
 PROCESS CA-84, IG-122
 PUT CA-87, IG-127
 PUTDE CA-89, IG-130
 PUTUP CA-91, IG-133
 RELEASE CA-93, IG-136
 syntax summary CA-95,
 IG-138
 verifying CA-152, IG-219
 index block header IM-6
 indexed access control block
 (IACB) IM-5
 installation CA-186, IG-221
 languages supporting CA-12,
 IG-4
 loading CA-15, IG-204
 loading file CA-33, IG-61
 locking, record CA-144
 master control block
 (MCB) IM-5
 overview SG-11
 performance CA-142, IG-205
 preparing application program
 \$JOBUTIL procedure CA-196,
 IG-230
 link-edit control
 file CA-197, IG-231
 overview CA-195, IG-229
 problem determination IM-57
 queue description IM-8
 record deletion
 processing CA-181, IM-21
 record insertion
 processing CA-175, IM-15
 record/request processing IM-2
 requesting records CA-15
 resource contention CA-144,
 IG-210
 return codes
 DELETE CA-68, IG-102
 DISCONN CA-69, IG-104
 ENDSEQ CA-71, IG-106
 EXTRACT CA-74, IG-109
 GET CA-77, IG-113
 GETSEQ CA-80, IG-116
 LOAD CA-83, IG-120
 PROCESS CA-86, IG-125
 PUT CA-88, IG-128
 PUTDE CA-89, IG-131
 PUTUP CA-92, IG-134
 RELEASE CA-93, IG-136
 summary CA-96, IG-139,
 MC-241
 Secondary Index
 Directory IM-29
 secondary indexes IM-25
 services CA-13
 storage allocation IM-1
 storage layout CA-15, IG-204
 storage requirements CA-141,
 IG-203

- tailoring
 - \$IAMUT1 buffer size CA-193
 - buffer/control block allocations CA-190
 - central buffer size CA-191
 - Indexed Access Method buffer size IG-44
 - using alternate allocations CA-189
 - using read-only \$IAM CA-194
 - volume requirements SG-178
 - with data paging IM-33
- indexed data set
 - See indexed file
- indexed file
 - building file CA-31
 - control block structure CA-155
 - design considerations CA-20
 - designing file CA-23
 - format
 - blocks CA-157, IG-24
 - clusters CA-168, IG-30
 - data blocks CA-160, IG-26
 - file control blocks CA-159, IG-40
 - free blocks CA-170, IG-27
 - free pool CA-172, IG-39
 - free records CA-168, IG-27
 - free space CA-168, IG-27
 - higher-level index
 - block CA-166, IG-37
 - index CA-161, IG-30
 - index blocks CA-159, IG-30
 - index example CA-167, IG-37
 - last cluster CA-171, IG-33
 - overview CA-154, IG-17
 - PIXB CA-161, IG-32
 - reserve blocks CA-170, IG-32
 - reserve index
 - entries CA-170, IG-36
 - sequential chaining CA-161, IG-28
 - SIXB CA-164, IG-35
 - loading with \$IAMUT1 IG-62
 - maintenance
 - backup and recovery CA-60, IG-94
 - deleting file CA-62, IG-96
 - dumping file CA-62, IG-96
 - recovery CA-61, IG-95
 - reorganizing file CA-61, IG-95
 - overview CA-18
 - processing
 - accessing file CA-55, IG-89
 - connecting file CA-54, IG-87
 - deleting records CA-59, IG-93
 - direct reading CA-55, IG-89
 - direct updating CA-56, IG-90
 - disconnecting file CA-54, IG-87
 - extracting file
 - attributes CA-59, IG-93
 - inserting records CA-59, IG-93
 - loading sequential file IG-65
 - sequential processing summary CA-58, IG-92
 - sequential reading CA-57, IG-90
 - sequential updating CA-57, IG-91
 - setting up
 - building file CA-31
 - defining file
 - structure IG-19
 - defining the key CA-24, IG-18
 - designing file
 - structure CA-25
 - loading base records CA-33, IG-61
 - loading from sequential file CA-36, IG-64
 - preparing input CA-35, IG-63
 - selecting block size CA-23
 - setting file
 - structure CA-31
 - structured and dynamic files CA-20, IG-40
 - indexing, address feature LR-7
 - indicator, auto-update IG-69
 - Information Management System/Virtual Storage (IMS/VIS) SI-1
 - INIT-SELF command SN-136
 - initial program load
 - See IPL (initial program load)
 - initialization
 - automatic restart SG-174
 - disk (4962) OR-131
 - disk (4963) OR-136
 - disk and diskette SG-204
 - disk or diskette, \$INITDSK utility OR-346
 - diskette
 - (4964,4965,4966) OR-122
 - graphics data base CT-198
 - IPL SG-204
 - module overview ID-61
 - modules ID-9
 - nucleus ID-5
 - Series1 to Series/1 CT-103
 - stand-alone dump
 - diskette OR-120
 - tape SG-205
 - terminal SG-213
 - initializing IPL text OR-358
 - initializing volumes OR-357
 - INITMODS table ID-9
 - INITTASK initialization task ID-5
 - input buffer, Multiple Terminal Manager
 - contents during I/O operations CA-220
 - description CA-219
 - internals M-23, M-37
 - input data parsing ID-240
 - input data set SM-9
 - input error function, Remote Management Utility ID-208
 - input exit
 - example SM-146
 - rules for passing records SM-139
 - uses SM-33
 - writing program with SM-137
 - input operations, GPIB ID-124
 - input record as data field
 - example SM-3
 - how to specify SM-125

required entries in field
 statement SM-109
 restriction with summary
 fields SM-126
 rules for specifying SM-125
 use with variable input
 records SM-125
INPUT subcommand
 (\$EDIT1/N) OR-222
 input/output control block
 See IOCB instruction
 input, terminal I/O LR-32
 insert secondary index directory
 entry IG-155
 inserting block of text OR-288
 inserting elements
 (graphics) CT-167
 inserting line of text OR-286
 inserting member
 (graphics) CT-177
 installing \$RJESNA SN-221
 installing the system SG-27
 instruction and statements over-
 view LR-13
 instruction definition and check-
 ing, \$EDXASM ID-267
 instruction expansions, 370 Chan-
 nel Attach T-29
 instruction format, Event Driven
 Language ID-55, LR-3
 instruction locations, directory
 of LR-34
 instruction processing,
 SNA SN-141
 instruction sequence,
 examples SN-13
 integer and logical
 instructions LR-14
 interactive debugging OR-142,
 SG-65
 intercepting class
 interrupts OR-491
 interchanging data SG-23
 interface layers SI-4
 interface routines,
 supervisor ID-50
 interface, \$JOBUTIL
 (\$EDXLINK) PP-65
 interface, \$JOBUTIL interactive
 (\$EDXLINK) PP-67
 interface, \$JOBUTIL
 non-interactive (\$EDXLINK) PP-66
 interface, \$L interactive (\$EDX-
 LINK) PP-62
 interface, \$L non-interactive
 (\$EDXLINK) PP-61
 interrupt information byte
 (IIB) CA-220
 interrupt keys
 attention key CT-20
 enter key CT-21
 program function (PF)
 keys CT-20
 interrupt line (Channel
 Attach) CT-466
 interrupt servicing
 EXIO ID-156
 GPIB ID-127
 immediate action
 routines ID-37
 Series/1 to Series/1 ID-132
 terminal I/O CT-20, ID-106
 interrupt types, PSW MC-282
INTIME instruction
 coding description LR-207
 overview LR-32, SG-42
 introduction to EDL LR-1
 INVITE parameter, NETPUT SN-30
 invoke your operator command
 (\$U) OR-37
 invoking \$GPIBUT1 OR-292
 invoking directory
 subcommands IG-150
 invoking operator commands OR-22
 invoking session manager OR-46
 invoking Sort/Merge, methods
 with \$JOBUTIL utility SM-42
 with LOAD SM-44
 invoking user written operator
 command SG-61
 invoking utilities SG-62
IOCB instruction
 BUFFER parameter LR-209
 coding description LR-209
 direct I/O LR-209
 function SG-79
 overview LR-30
 structure CT-35
 terminal I/O
 instruction LR-209
 use in accessing virtual termi-
 nals CT-205
 use in defining logical
 screen CT-33
 using PRINTTEXT LR-247
 using READTEXT LR-300
IODEF statement
 coding description SG-361
 function SG-59, SG-360
 SPECPI process interrupt user
 routine SG-365
IOR instruction
 coding description LR-214
 overview LR-14
IOSPTBL (spool I/O control block)
 description ID-151
 listing ID-371
IP command (\$S1S1UT1) CT-103
IPL (initial program load)
 \$S1S1UT1 utility CT-86
 automatic
 initialization SG-174
 diagnosing a problem SG-270
 internals ID-5
ISAPPID parameter, NETHOST SN-17
ISMODE parameter, NETHOST SN-17
ISPASWD parameter, NETHOST SN-18
ISQUEUE parameter, NETHOST SN-18
ISRQID parameter, NETHOST SN-18
ISUSFLD parameter, NETHOST SN-18
IV command (\$INITDSK) OR-357

J

JOB command (\$JOBUTIL) OR-384
 job control block
 spool active job (SPA)
 description ID-147
 listing ID-371
 spool job (SPJ)
 description ID-146
 listing ID-372
 job control statement OR-384
 job stream control records SN-206

job stream processor,
 \$JOBUTIL OR-376
 job streams, submitting SN-214
 JOURNAL command (\$RJESNA) SN-212
 journal data set SN-212
 journaling messages SN-212
 JP command (\$PDS) SG-329
 JP subcommand (\$DICOMP) CT-177
 JR subcommand (\$DICOMP) CT-178
 JUMP command (\$JOBUTIL) OR-385
 jump to address (graphics) CT-177
 jump to reference
 (graphics) CT-178
 jumping to batch job stream
 label OR-385

K

keep/drop option
 specifying SM-64
 when to use SM-65
 keeping spool job CA-427
 key
 See program function (PF) keys
 key, primary IG-18
 key, secondary IG-67
 Keyboard define utility for 4978,
 \$TERMUT2 OR-468
 KEYS command (\$IMAGE) OR-331
 keyword operand definition LR-5

L

LA command (\$DISKUT1) OR-168
 LA command (\$DIUTIL) CT-199
 LA command (\$TERMUT1) OR-465
 label
 field LR-3
 syntax description LR-4
 LABEL command (\$JOBUTIL) OR-386
 LABELS subroutine, \$EDXASM ID-264
 labels, tape
 See tape labels
 LACTS command (\$DISKUT1) OR-168
 LAD command (\$DISKUT1) OR-178
 language control data set,
 \$EDXASM ID-243, PP-22
 language products SG-3
 LAST parameter, NETPUT SN-30
 LASTQ instruction
 coding description LR-216
 overview LR-24, SG-42
 LAV command (\$DISKUT1) OR-179
 layers, terminal I/O ID-97
 LB function (\$PDS) SG-326
 LB subcommand (\$DICOMP) CT-178
 LC command (\$TERMUT2) OR-476
 LD command (\$DISKUT1) OR-174
 LD command (\$IOTEST) OR-374
 LDCB command (\$GPIBUT1) CT-137
 LDCTS command (\$DISKUT1) OR-168
 LE (less than or equal) LR-21
 LE command (\$IAMUT1) IG-156
 level status block
 See LSB (level status block)
 LH command (\$DIUTIL) CT-199
 LI command (\$TERMUT2) OR-477
 LI function (\$PDS) SG-327

LI subcommand (\$DICOMP) CT-178
 library
 definition SG-48
 directory ID-77
 licensed programs
 descriptions SG-5
 Event Driven Language SG-26
 FORTRAN SG-25
 Macro Assembler Language SG-26
 line command (\$FSEDIT) OR-285
 line commands, \$FSEDIT OR-280
 line connection types
 (SDLC) SN-105
 line continuation, source LR-4
 line editing, \$EDIT1/N OR-235
 description OR-237
 line pointer reposition
 See moving line pointer
 LINK control statement
 (\$EDXLINK) PP-54
 link map (\$EDXLINK) PP-70
 link-editing \$NETCMD SN-12
 LINK, Multiple Terminal Manager
 CALL
 coding description CA-260
 internals M-11
 overview of CA-206
 linkage editor
 See \$LINK
 linkage editor overlays
 See overlay support
 linkages AM-9
 linking decompaction
 routines SN-224
 LINKON, Multiple Terminal Manager
 CALL
 coding description CA-262
 internals M-11
 overview of CA-206
 LIST command (\$DEBUG) OR-154
 LIST command (\$EDIT1/N) OR-206
 with \$EDIT1/N
 READ OR-207
 LIST command (\$EDXASM) PP-13
 LIST command (\$EDXLIST) PP-25
 list partitions and
 programs OR-24
 LIST subcommand (\$EDIT1/N) OR-224
 listing breakpoints and trace
 ranges OR-151
 listing characters
 (graphics) CT-178
 listing control instructions
 EJECT LR-110
 overview LR-19
 PRINT LR-244
 SPACE LR-322
 TITLE LR-380
 listing data members OR-174
 listing data members in CTS
 mode OR-168
 listing data set
 with \$DISKUT2 OR-185
 with \$DIUTIL CT-202
 with \$EDIT1/N OR-224
 with \$FSEDIT OR-270
 listing data set status
 (graphics) CT-201
 listing date/time CT-183, OR-44
 listing device control block
 (GPIB) CT-137
 listing devices OR-374
 listing error specification PP-13

listing hardware configuration OR-368
 listing insert mask OR-289
 listing members CT-168, CT-199, OR-175
 listing members, all
 with \$DISKUT1 OR-168
 with \$DIUTIL CT-199
 listing processor program, \$EDX-LIST PP-24
 listing program function key codes OR-410
 listing program function keys OR-331
 listing program members OR-168
 listing secondary index directory IG-156
 listing status of all tasks OR-162
 listing storage OR-30
 listing supervisor configuration OR-375
 listing terminal names/types/addresses OR-465
 listing variables CT-183
 listing volumes on disk or diskette OR-362
 LISTP command (\$DISKUT1) OR-168
 LISTT command (\$DISKUT1) OR-168
 LL command (\$DISKUT1) OR-175
 LL command (\$DISKUT2) OR-186
 LO command (\$IAMUT1) CA-110, IG-162
 LOAD Indexed Access Method CALL
 description CA-81, IG-118
 return codes CA-83, IG-120
 LOAD instruction
 coding description LR-217
 example SM-146
 overview LR-28
 preparing a job using SM-131
 required parameters SM-132
 return codes LR-221, MC-244
 return codes from Sort/Merge SM-135
 used with overlays SG-40
 when to use SM-131
 with automatic initialization SG-174
 load mode IG-61
 load point defined LR-26
 loading \$SNA SI-3
 loading control store OR-476
 loading image store OR-477
 loading indexed file CA-110, IG-162
 loading overlay programs ID-19
 loading overlay segments ID-20
 loading program
 \$L operator command OR-32
 automatic initialization SG-174
 EXEC \$JOBUTIL command OR-383
 internals ID-16
 loading secondary index IG-75
 loading SNA support SN-100
 LOCAL ATTNLIST LR-50
 LOCATE command (\$FSEDIT) OR-277
 location dictionary, \$EDXASM ID-277
 locking queues AM-10, IM-11
 locks, Indexed Access Method
 block and record AM-10, IM-11
 conditional processing IG-88
 description CA-54, IG-87
 file AM-12, IM-13
 overview CA-3, IG-2
 LOG command (\$JOBUTIL) OR-387
 log data set for I/O errors OR-400
 LOG macro SG-263
 log statement
 default if not written SM-57
 how to write SM-58
 purpose SM-57
 relative position SM-57
 logging device for messages
 changing SM-59
 defaults SM-59
 logging on to host (\$RJESNA) SN-207
 logical end-of-file on disk SG-354
 logical screen, terminal support
 examples CT-33, CT-34
 using IOCB and ENQT to define CT-33
 using TERMINAL to define CT-32
 logical unit, defining SN-114
 logon menu, session manager OR-46
 low storage
 during IPL ID-6
 during program load ID-16
 LP command (\$DISKUT1) OR-176
 LP command (\$DISKUT2) OR-181
 LPCTS command (\$DISKUT1) OR-168
 LR function (\$PDS) SG-331
 LR subcommand (\$DICOMP) CT-179
 LS command (\$DISKUT1) OR-177
 LS command (\$IOTEST) OR-375
 LSB (level status block) ID-42
 LT (less than) LR-21
 LU command (\$DISKUT2) OR-185
 LU parameter
 NETCTL SN-48
 NETGET SN-37
 NETINIT SN-19
 NETPUT SN-30
 NETTERM SN-59
 SNALU SN-114
 LUE control block, LUE SI-342
 LUL control block, LUL SI-342, SI-352
 LUS, system services control block SI-366, SI-367
 LUSTAT command SN-136

M

macro assembler internals ID-279
 macro library SG-7
 macro library/host SG-8
 magazine, diskette
 See 4966 Diskette Magazine Unit
 magnetic tape
 See tape
 mapping storage SG-104
 MASK line command (\$FSEDIT) OR-289
 master control block
 description IM-5
 printer spooling description ID-144
 listing ID-373
 See also MCB (master control block)

Mathematical and Functional
 Subroutine Library SG-8
 MCB (master control block)
 description AM-5, IM-5
 equates AM-51
 MCB (member control block) SG-333
 MD command (\$DIUTIL) CT-200
 member area SG-324
 member control block (MCB) SG-333
 MENU command (\$FSEDIT) OR-277
 MENU Multiple Terminal Manager
 CALL
 coding description CA-264
 overview of CA-206
 MENU, Multiple Terminal Manager
 CALL
 internals M-11
 menus
 option selection menu OR-52
 parameter selection menu OR-52
 primary option menu OR-55
 primary option menus OR-264
 See also session manager, menus
 See session manager, menus
 MERGE (merge)
 contrasted with sort SM-1
 example SM-16
 functions supported SM-13
 output record size,
 determining SM-242
 specifying in header SM-60
 MERGE option (\$FSEDIT) OR-271
 message acknowledgment SN-33
 message cancellation SN-52
 message completion criteria SN-39
 message flow, suspending SN-53
 Message Format Service
 (MFS) SN-161
 message length, specifying SN-30
 message numbers SN-24
 message output
 changing output device SM-57
 overriding header
 statement SM-60
 specifying SM-64
 with header statement SM-64
 with log statement SM-60
 message resynchronization
 data set SN-63
 data set content SN-64
 data set size SN-63
 disk considerations SN-65
 main storage
 considerations SN-65
 protocols for STSN SN-65
 message resynchronization, defi-
 nition of SN-7
 message sending utility,
 \$TERMUT3 OR-487
 message syntax MC-3
 message verification SN-52
 message verification,
 requesting SN-30
 message/module cross-reference,
 Sort/Merge SM-193
 message, PRINTEXT
 instruction LR-248
 messages
 compilation time
 (\$EDXASM) MC-169
 compilation time
 (\$S1ASM) MC-191
 definition SN-6
 examples with
 explanation SM-36
 execution time
 (alphabetized) MC-11
 execution time (with
 variables) MC-5
 list SM-249
 program status SM-247
 specification SM-247
 statistics SM-247
 system initialization SG-204
 warning message option SM-66
 messages, definition of SN-6
 MFA
 See Multifunction Attachment
 mixed precision
 combinations LR-15
 MM line command (\$FSEDIT) OR-290
 mode table, patching
 \$RJESNA SN-227
 mode, load IG-61
 mode, processing IG-61
 modified data
 reading from the 3101 CT-57
 reading from the 4978 CT-57
 3101 considerations CT-71
 3101 example CT-73
 modified data tags CT-71
 modify defaults, Remote Management
 Utility CT-320
 modifying character image
 screen OR-246
 modifying character string
 with \$EDIT1/N OR-214
 with \$FSEDIT OR-273
 modifying default storage allo-
 cation OR-190
 modifying existing data
 set OR-267
 modifying image/control
 store OR-468
 modifying storage or
 registers OR-157
 modifying the TERMINAL
 statement ID-116
 module descriptions
 GPIB ID-71
 Indexed Access Method
 common AM-23, IM-45
 data paging IM-53
 secondary index IM-55
 Multiple Terminal Manager
 component M-25
 link modules T-18
 overlay programs T-32
 program modules T-8
 utility M-29
 utility program
 modules T-26
 printer spooling ID-71
 Remote Management
 Utility ID-217
 Series/1 to Series/1 ID-72
 supervisor ID-57
 370 Channel Attach T-7, T-25
 module names, supervisor SG-465
 module removal, \$IAMSTGM IG-206
 MOVE instruction
 coding description LR-225
 overview LR-14
 MOVEA instruction
 coding description LR-229
 overview LR-14

MOVEBYTE subroutine,
 \$EDXASM ID-262
 moving block of text OR-290
 \$FSEDIT line command OR-290
 moving line pointer
 down OR-219
 to bottom OR-213
 to top OR-232
 up OR-233
 moving tape OR-449
 moving text
 \$EDIT1/N subcommand OR-225
 moving volumes on disk or
 diskette OR-404
 MP function (\$PDS) SG-327
 MP subcommand (\$DICOMP) CT-180
 MSGDATA parameter, NETINIT SN-20
 MT command (\$TAPEUT1) OR-449
 MTMFILES utility CA-319, M-29
 MTMPFMAP utility CA-315, M-29
 MTMPGMS utility CA-316, M-29
 MTMRPT utility CA-318, M-29
 MTMSCRNS utility CA-317, M-29
 MTMSTORE file CA-212, M-20
 MTMSTR volume CA-218
 MTMUPDT utility CA-312, M-29
 MTMUT2 utility CA-308, M-30
 multiconditional force
 data types allowed SM-112
 definition SM-119
 how to specify SM-119
 keep/drop from output
 record SM-64
 length, maximum
 in a field statement SM-112
 specifying in header SM-63
 total in include set SM-119
 number, maximum SM-119
 required entries in field
 statement SM-109
 rules for SM-115
 Multifunction Attachment
 attachment with 3101 Display
 Terminal CT-477
 changing random access
 memory SG-212
 considerations for attachment
 of devices SG-113
 defining SG-90
 description CT-9, CT-467
 module description ID-69
 use in binary synchronous
 communications CT-11
 use with EDX terminals CT-9
 multiple overlay areas,
 \$EDXASM ID-249
 Multiple Terminal Manager
 accessing the TEB CA-233
 adding a new device M-6
 application execution CA-217
 application
 organization CA-219
 application program design
 considerations
 COBOL CA-283
 EDL CA-281
 FILEIO CA-290
 FORTRAN CA-285
 PL/I CA-287
 spooling CA-289
 application program stub M-9
 application swap-out CA-218
 automatic OPEN/CLOSE CA-242
 CALL functions
 ACTION CA-225, M-10
 BEEP CA-229, M-9
 BLINK CA-231, M-10
 CDATA CA-233, M-9
 CHGPAN CA-236, M-9
 CYCLE CA-238, M-10
 FAN CA-240
 FILEIO CA-242, M-9
 FTAB CA-254, M-26
 GETCUR CA-258, M-10
 LINK CA-260, M-11
 LINKON CA-262, M-11
 MENU CA-264, M-11
 overview CA-205
 SETCUR CA-266, M-9
 SETFMT CA-268, M-10
 SETPAN CA-272, M-9
 WRITE CA-275, M-10
 CCB with printer spooling M-13
 component flow M-3
 components CA-205, M-3
 concepts and facilities CA-201
 creating additional
 copies CA-386
 data files
 MTMSTORE M-20
 MTMSTR volume M-20
 PRGRMS volume M-5
 data files/volumes
 considerations CA-290
 MTMSTORE file CA-212
 MTMSTR volume CA-218
 PRGRMS volume CA-212
 SCRNS volume CA-213
 SIGNONFL file CA-293
 TERMINAL file CA-212
 debugging applications CA-300
 direct file
 FCA format CA-244
 request types CA-247
 disk file support CA-208
 distribution CA-362
 ending application CA-304
 ending manager
 operation CA-304
 equate tables M-31
 examples
 dynamic screen
 creation CA-324
 dynamic screen
 modification CA-323
 file maintenance transaction
 application CA-327
 indexed file update CA-326
 sample signon
 program CA-352
 sample terminal I/O error
 program CA-352
 screen display
 techniques CA-325
 screen retrieval and
 display CA-321
 file control area (FCA) CA-243
 file management M-8
 file manager CA-208
 FILEIO direct file
 access CA-290
 FILEIO indexed file
 access CA-291
 hardware requirements CA-201
 indexed file
 considerations CA-291
 FCA format CA-245
 request types CA-246

input buffer address CA-219
 input buffer contents during
 I/O operations CA-220
 installation CA-369
 interrupt information byte
 (IIB) CA-220
 introduction CA-4
 link-edit considerations
 COBOL CA-298
 EDL CA-296
 FORTRAN CA-297
 PL/I CA-299
 module descriptions M-25
 output buffer CA-220
 overview SG-12
 printer spooling, with M-13
 program management M-3
 program manager CA-205
 programming considerations
 COBOL CA-283
 EDL CA-281
 FORTRAN CA-285
 PL/I CA-287
 return codes
 FILEIO CA-247
 FTAB CA-255
 SETFMT CA-270
 SETPAN CA-273
 summary CA-278, MC-246
 screen formats
 SCRADD CA-313
 SCRNFIO1 CA-319
 SCRNFIO2 CA-320
 SCRNIPL CA-214
 SCRNMENU CA-216
 SCRNM1 CA-308
 SCRNM2 CA-309
 SCRNM3 CA-310
 SCRNM4 CA-311
 SCRNPMP CA-315
 SCRNPGM CA-314
 SCRNSIGN CA-215
 SCRNSREP CA-317
 SCRNPDT CA-312
 screen image manager M-7
 signing on and off CA-212
 SIGNOFF CA-295
 SIGNON CA-293
 software requirements CA-204
 spooling considerations CA-289
 storage environment M-11
 storage requirements SG-222
 summary CA-278
 swap-out data set CA-218
 system generation consider-
 ations CA-367, SG-176
 tailoring CA-376
 terminal environment block
 (TEB) CA-220, M-17
 terminal manager CA-207
 terminal servers M-6
 terminal states M-39
 utilities
 \$MTMUT1 CA-307
 \$RECON CA-306
 DISCONNECT CA-306
 MTMFILES CA-319
 MTMPFMAP CA-315
 MTMPGMS CA-316
 MTMRPT CA-318
 MTMSCRNS CA-317
 MTMUPTD CA-312
 MTMUT2 CA-308
 RECONNECT CA-306

3101 considerations CA-202
 MULTIPLY instruction
 coding description LR-230
 overview LR-14
 precision table LR-231
 multiprogramming
 automatic
 initialization SG-174
 design feature SG-16
 multitasking SG-16

N

NBLINK subroutine CA-476
 NCP major node definition
 operands SN-313
 sample SN-318
 Series/1 related
 parameters SN-313
 NE (not equal) LR-21
 negative acknowledgments, respond-
 ing to SN-33
 NETBIND return codes MC-260,
 SN-94
 NETCLOSE MC-261
 NETCLOSE return codes MC-261,
 SN-93
 NETOPEN MC-262
 NETCTL instruction
 coding examples SN-54
 description SN-47
 overview SN-9
 return codes MC-251, SN-55
 summary of protocols SN-50
 syntax SN-47
 types of control
 operations SN-49
 NETGET instruction
 coding example SN-43
 description SN-37
 overview SN-8
 return codes MC-253, SN-44
 summary of protocols SN-38
 syntax SN-37
 NETHOST instruction
 description SN-17
 syntax SN-17
 NETINIT instruction
 coding examples SN-17
 description SN-19
 overview SN-6
 return codes MC-256, SN-28
 return codes from STSN process-
 ing SN-85
 summary of protocols SN-23
 syntax SN-19
 NETOPEN return codes MC-262,
 SN-89
 NETRECV MC-263
 NETPUT instruction
 coding examples SN-34
 description SN-29
 overview SN-6
 return codes MC-257, SN-36
 summary of protocols SN-31
 syntax SN-29
 NETRECV return codes MC-263,
 SN-90
 NETSEND MC-264
 NETSEND return codes MC-264,
 SN-91

NETUBND MC-267
 NETTERM instruction
 coding example SN-61
 description SN-59
 overview SN-10
 return codes MC-258, SN-62
 syntax SN-59
 NETUBND return codes MC-267,
 SN-94
 session termination MC-268
 network activation
 checklist SN-311
 network control program
 (NCP) SI-1
 newline subroutine ID-106
 NEXTQ instruction
 coding description LR-233
 overview LR-24
 timers SG-42
 no data record, PASSTHRU function
 of \$RMU CT-293
 NOFLOAT module description ID-66
 NOLIST command (\$EDXASM) PP-13
 NOMSG command (\$JOBUTIL) OR-388
 non-compressed byte string CT-58
 non-labeled tapes
 description SG-299
 layout SG-299
 processing SG-300
 NORING parameter, SNAPU SN-105
 normal control field
 contrast with opposite control
 field SM-116
 data types allowed SM-112
 dropping SM-64
 example 1-13 SM-29
 how to specify SM-116
 keep/drop from output
 record SM-64
 length, maximum
 in a field statement SM-112
 specifying in header SM-64
 required entries in field
 statement SM-109
 rules SM-121
 specifying in header SM-64
 when to drop SM-65
 NOTE instruction
 coding description LR-236
 function SG-57
 overview LR-16
 Version 3 conversion consider-
 ations SG-435
 notify of an event
 See POST instruction
 NP command (\$IAMUT1) IG-166
 NQ command (\$COPYUT1) OR-113
 NSPE command SN-136
 nucleus initialization ID-5
 NULL character, defining OR-332
 NULL command (\$IMAGE) OR-332

0

object data set for \$EDXASM PP-11
 object module record
 format ID-381
 object text elements,
 \$EDXASM ID-237
 obtaining a stand-alone
 dump SG-240

OFF function, CONTROL
 instruction LR-65
 OLE (operand list element) \$EDXASM
 format ID-239
 in instruction parsing ID-242
 used in \$IDEF ID-267
 omit set
 definition SM-93
 how to specify SM-95
 omit statement
 how to specify SM-95
 purpose SM-93
 op (operation field) LR-3
 OPCHECK subroutine,
 \$EDXASM ID-258
 opening a data set
 on disk or diskette ID-90
 on tape ID-94
 opening EXIO device,
 EXOPEN ID-155
 opening member, \$PDS
 utility SG-334
 OPENT subroutine CA-477
 operands
 defined LR-3
 keyword LR-5
 parameter naming (Px) LR-8
 operating environment
 minimum system
 configuration SG-23
 required for program prepara-
 tion SG-23
 operation codes
 description ID-56
 emulator command table
 listing ID-323
 in \$EDXASM compiler ID-242
 operations using index
 registers LR-15
 operator command overview
 (\$RJESNA) OR-421, SN-210
 operator command/utility selection
 guide OR-9
 operator commands OR-21, SG-61
 operator commands overview OR-2
 operator termination
 (\$EDXLINK) PP-69
 operators, arithmetic LR-4
 opposite control field
 contrast with normal control
 field SM-116
 data types allowed SM-112
 dropping when not
 ascending SM-65
 examples SM-18
 how to specify SM-116
 keep/drop from output record
 default drop with variable
 records SM-65
 example SM-20
 specifying in header SM-64
 when to drop SM-65
 length, maximum
 in a field statement SM-112
 specifying in header SM-64
 total in include set SM-116
 number, maximum SM-116
 required entries in field
 statement SM-109
 rules for (see also control
 field) SM-116
 option selection menu OR-52
 optional features support LR-13
 orderly session termination SN-60

OTE \$EDXASM instruction ID-253
 output buffer, Multiple Terminal Manager
 contents during I/O
 operations CA-220
 description CA-220
 internals M-23, M-37
 output considerations for terminal support
 end of forms for printers CT-24
 forms control CT-23
 output line buffering CT-23
 OUTPUT control record (\$LINK) PP-89
 output data set SM-9
 output exit
 example SM-146
 returning control SM-142
 rules for passing output records SM-142
 uses SM-36
 writing program with SM-139
 output operations, GPIB ID-125
 output record size, determining SM-241
 overflow indicator
 example SM-29
 how to specify SM-127
 purpose SM-127
 rules SM-127
 overflow, arithmetic LR-15
 OVERLAY control statement (\$EDX-LINK) PP-55
 overlay function process table ID-192
 OVERLAY option (\$EDXASM) PP-14
 overlay support
 overlay program SG-37, T-30
 loading ID-19
 subroutines, \$EDXASM ID-257
 supervisor initialization routines SG-210
 overlay segment SG-37
 considerations PP-78
 loading ID-20
 storage layout PP-78
 overlay table, Remote Management Utility ID-192
 overlay work area, \$EDXASM ID-247

P

pacing SN-131
 pacing response SN-132
 pad characters, specifying SN-105
 PAD parameter, SNAPU SN-105
 page eject SG-61
 page table, data paging IM-36
 parameter equate tables LR-431
 parameter list SM-132
 parameter naming operands (Px=) SN-11
 parameter naming operands in instruction format LR-8, SN-11
 parsing, input data (\$EDXASM) ID-240
 partial messages, sending SN-35
 partition assignment, changing CP command (\$GPIBUT1) CT-135
 partitioned data sets SG-321

partitions SG-43
 pass control to log routines SG-263, SG-265
 passing parameters using \$JOBUTIL OR-389
 PASSTHRU function
 communications flow CT-284
 conducting a session CT-286
 establishing a session CT-281
 example of passing parameters CT-256
 internals ID-205
 logic flow diagram CT-288
 overview CT-281
 programming considerations CT-294
 record blocking CT-294
 sample program CT-306
 types of records CT-289
 virtual terminals CT-294
 with \$DEBUG CT-314
 PASSTHRU record, Remote Management Utility CT-253
 PATCH command (\$DEBUG) OR-157
 patch, Remote Management Utility defaults CT-320
 patching \$RJESNA mode table SN-227
 patching disk/diskette OR-181
 patching storage OR-33, SG-61
 patching storage or registers OR-157
 PAUSE command (\$JOBUTIL) OR-390
 PC function (\$PDS) SG-329
 PC subcommand (\$DICOMP) CT-181
 PD command (\$IOTEST) OR-370
 performance (Sort/Merge) SM-155
 performance aids
 supervisor initialization routines SG-210
 using 1024-byte sectors SG-208
 performance volume defined SG-96
 description SG-208
 performing a customized system generation SG-194
 PF keys
 See program function (PF) keys
 PG command (\$IAMUT1) IG-167
 GPIB command (\$GPIBUT1) CT-142
 phase execution, \$SIASM ID-281
 physical unit services control block, PUS SI-371, SI-372
 physical unit, defining SN-104
 PI
 See process interrupt
 PID program directory SG-27
 PIXB (primary-level index block) CA-161, IG-32
 PL subcommand (\$DICOMP) CT-181, CT-182
 realtime data member CT-184
 PL/I
 estimating storage requirements SG-228
 execution requirements SG-25
 link editing SG-72
 overview SG-10
 program preparation requirements SG-25
 use with Multiple Terminal Manager CA-287
 PLC, program list control block SI-369, SI-370

plot control block
 (graphics) CT-156
 plot curve data member, \$PDS utility SG-325
 PLOTCB control block CT-156
 PLOTGIN instruction
 coding description CT-155
 overview CT-155
 PLU bracket initiation SN-128
 POINT instruction
 coding description LR-238
 function SG-57
 overview LR-16
 polling and selection, BSC CT-215
 post codes
 BIND event MC-218, SN-94
 channel attach CT-438, MC-219
 See also completion codes and return codes
 tape CONTROL LR-67, MC-222
 tape READ LR-296, MC-222
 tape WRITE LR-397, MC-222
 POST command (\$DEBUG) OR-159
 POST instruction
 coding description LR-240
 internals ID-48
 overview LR-28, SG-31
 POST supervisor function ID-36
 posting GPIB operation
 complete CT-142
 power outage, restoring
 after SG-174
 PP command (\$IAMUT1) IG-168
 PR command (\$CHANUT1) CT-409
 PR command (\$DICOMP) CT-168
 precision
 combinations allowed LR-15
 floating-point
 arithmetic LR-16
 integer and logical LR-14
 precision table LR-15
 PREPARE IDCBC command LR-201
 PRGRMS volume CA-212
 PRGRMS volume, Multiple Terminal Manager M-5
 primary OR-264
 primary commands, \$FSEDIT OR-272
 primary control statement data
 set PP-48
 primary option menu,
 \$FSEDIT OR-264
 primary option menu, session manager OR-55
 primary-level index block (PIXB) CA-161, IG-32
 PRINDATE instruction
 coding description LR-242
 function SG-79
 overview LR-30, SG-42
 PRINT statement
 coding description LR-244
 overview LR-19
 print utilities, remote job
 entry CT-343
 printer function, \$RJESNA SN-204
 printer spooling
 advantages of CA-3
 automatic writer start CA-441
 capacity parameters
 changing CA-437
 defaults CA-439
 cold starting CA-438
 components
 operator command CA-405

spool manager CA-400
 spool utility CA-404
 spool writers CA-402
 concepts and facilities CA-393
 considerations for
 Multiple Terminal Manager user CA-289
 PL/I user CA-408
 control blocks
 descriptions ID-143
 listings ID-371
 controlling from a program
 controlling spool job processing CA-449
 finding if spooling active CA-454
 preventing spooled output CA-447
 separating spooled output CA-449
 summary CA-447
 cross-reference CA-409
 data flow CA-397
 functions summary CA-393
 installing and optimizing CA-455
 internals
 \$SPOOL program ID-135
 control block
 listings ID-371
 control blocks ID-143
 introduction CA-396
 language considerations CA-408
 Multiple Terminal Manager M-13
 operator functions
 change max active spool jobs CA-441
 change max spool jobs CA-441
 change spool data set CA-441
 change spool devices CA-441
 delete generic spool jobs CA-420
 delete spool jobs CA-418
 display spool status CA-421
 hold spool job(s) CA-426
 invoking CA-412
 keep spool job CA-427
 release spool job(s) CA-429
 restart spool writer CA-431
 start spool writer CA-435
 stop spool facility CA-430
 stop spool writer CA-433
 overview SG-18
 requirements to use CA-397
 restarting CA-438
 return codes MC-249
 starting CA-401
 storage requirements SG-219
 supervisor interface module ID-139
 system generation CA-397
 system generation considerations SG-176
 term definitions CA-395
 terminating CA-402
 tips and techniques CA-406
 PRINTEXT instruction
 coding description LR-247
 function SG-79
 overview LR-30
 return codes LR-253
 use in terminal support

- changing individual fields CT-40
- using on 3101 terminals CT-44
- PRINTIME instruction
 - coding description LR-259
 - function SG-79
 - overview LR-30
- printing BSC trace files CT-237
- printing graphics member CT-168
- PRINTNUM instruction
 - coding description LR-261
 - function SG-79
 - overview LR-30
- PRINTON command (\$RJESNA) SN-213
- PRINTON command (\$RJE2780/\$RJE3780) CT-340
- priorities, \$RJESNA workstation SN-205
- priorities, SNA task SN-11
- priority
 - assigned to tasks SG-30
 - design feature SG-16
 - illustrated SG-34
 - program LR-267
 - SNA task SN-11
 - task LR-335
 - task internals ID-30
 - work station (\$RJESNA) SN-205
- PRNT command (\$IMAGE) OR-334
- problem determination
 - access method AM-31, IM-57
 - binary synchronous communications CT-235
 - BSC utilities for CT-10
 - Remote Management Utility ID-215
 - Sort/Merge SM-179
 - terminal I/O ID-112
- PROC command (\$JOBUTIL) OR-392
- procedures, session manager
 - See session manager
- PROCESS Indexed Access Method CALL
 - description CA-84, IG-122
 - return codes CA-86, IG-125
- process interrupt
 - control block (SBIOCB) ID-159
 - description SG-357
 - IODEF statement SG-361
 - IOTEST command OR-368
 - supported by sensor I/O SG-20
 - user routine (SPECPI) SG-364
- processing compacted data SN-223
- processing compiler output with \$LINK or \$UPDATE PP-7
- processing mode IG-61
- processor status word (PSW) MC-282, SG-275
- processor-to-processor communications
 - transmission protocol CT-81
 - use in Series/1 to Series/1 Attachment CT-87
- PROGEQU (program header equate table) LR-431
- program assembly or compilation PP-2
- program check message, formats MC-279, SG-273
- program control LR-20
- Program Design Language (PDL)
 - listing conventions SI-6
 - overview SI-6
 - verbs SI-7

- program design language, SDLC layer modules
 - CDCABNO SI-244
 - CDCADDTO SI-246
 - CDCATTN SI-247
 - CDCBLDDC SI-248
 - CDCCL500 SI-250
 - CDCCMDPR SI-251
 - CDCCMDR SI-254
 - CDCCNTBU SI-255
 - CDCCNTRL SI-256
 - CDCDISAB SI-257
 - CDCENABL SI-259
 - CDCIHT SI-260
 - CDCISXMT SI-267
 - CDCNOI SI-271
 - CDCOIO SI-272
 - CDCPERME SI-274
 - CDCPURGE SI-275
 - CDCRCSS SI-276
 - CDCSETRE SI-278
- program design language, Series/1 SNA base layer modules
 - CDDACTLU SI-144
 - CDDBETBR SI-145
 - CDDBSMGR SI-146
 - CDDCLEAR SI-147
 - CDDCOMSC SI-148
 - CDDCPMGS SI-150
 - CDDCPMGX SI-151
 - CDDCP2S SI-152
 - CDDCTDLT SI-153
 - CDDCTGEN SI-154
 - CDDCTPRC SI-156
 - CDDDRQRV SI-157
 - CDDEXRCV SI-158
 - CDDFISC SI-159
 - CDDGTBFS SI-161
 - CDDHXMGR SI-162
 - CDDINMSG SI-165
 - CDDISEG SI-170
 - CDDLUPRG SI-172
 - CDDMRSCV SI-175
 - CDDNTDAC SI-176
 - CDDNTRCV SI-178
 - CDDNTSND SI-180
 - CDDPCLUR SI-186
 - CDDPCNGS SI-190
 - CDDPCPUR SI-191
 - CDDPC2OT SI-194
 - CDDPSV00 SI-195
 - CDDPUSIN SI-198
 - CDDQCHK SI-199
 - CDDQOFF SI-200
 - CDDQON SI-201
 - CDDRECVY SI-202
 - CDDRELBL SI-203
 - CDDRELBS SI-204
 - CDDRQCV SI-174
 - CDDRQRCV SI-205
 - CDDRQRSZ SI-211
 - CDDRQSNN SI-212
 - CDDRQSNX SI-219
 - CDDRSRCV SI-220
 - CDDRSSNN SI-222
 - CDDRSSNX SI-223
 - CDDSELB SI-224
 - CDDSESEL SI-225
 - CDDSTPLK SI-227
 - CDDSVCCR SI-228
 - CDDZPNMS SI-229
 - NETBIND SI-134
 - NETCLOSE SI-140
 - NETOPEN SI-129

NETUBND SI-139
 program design language, Series/1
 SNA top layer modules
 CDDFANE SI-98
 CDEATTN SI-24
 CDEBUFF SI-50
 CDECD SI-25
 CDECHKA SI-26
 CDECHKRU SI-27
 CDECHK00 SI-33
 CDECLEAN SI-34
 CDECTLOA SI-36
 CDECTLOB SI-38
 CDECTLOC SI-40
 CDECTLOD SI-41
 CDECTLOO SI-42
 CDECTLO1 SI-44
 CDEDEQ SI-46
 CDEDEQAB SI-47
 CDEENQ SI-48
 CDEEXIT SI-49
 CDEGBUFF SI-51
 CDEGET00 SI-52
 CDEGET01 SI-53
 CDEINITA SI-59
 CDEINITO SI-61
 CDEINIT1 SI-62
 CDEPOST SI-64
 CDEPUT0A SI-65
 CDEPUT00 SI-69
 CDEPUT01 SI-72
 CDERCVRU SI-75
 CDEREJ SI-76
 CDERSP SI-77
 CDERSYNA SI-78
 CDERSYNC SI-79
 CDESNDRU SI-81
 CDESTSN SI-83
 CDESTSNA SI-85
 CDETERMA SI-88
 CDETERM0 SI-91
 CDETERM1 SI-95
 program end record, PASSTHRU function of \$RMU CT-293
 program execution via Remote Management Utility
 EXEC function CT-264
 PASSTHRU function CT-281
 SHUTDOWN function CT-276
 program function (PF) keys
 internals ID-102
 listing 4978, \$PFMAP utility OR-410
 PF6 with spool writer active CA-407
 use in terminal support CT-20
 use with attention lists CT-21
 when using \$FONT OR-252
 when using \$FSEDIT OR-262
 when using \$IMAGE OR-331
 when using session manager OR-70
 Program Function key record, PASSTHRU function of \$RMU CT-289
 program header
 description ID-21
 listing ID-365
 program identifier, \$JOBUTIL command OR-393
 PROGRAM instruction
 coding description LR-266
 overview LR-28, SG-31
 program header generation ID-21

program library update
 See \$UPDATE utility
 program library update (see \$UPDATE)
 program list control block, PLC SI-369, SI-370
 program loading
 See LOAD instruction
 program module sectioning functions LR-20
 program organization, Sort/Merge SM-180
 program preparation
 \$EDXASM PP-11
 \$SIASM PP-27
 example SG-383
 host assembler PP-39
 on System/370 PP-4, SG-7
 overview PP-1, SG-6, SG-21
 Remote Mangement Utility ID-210
 summarized SG-66
 utilities OR-56, PP-1
 program sequencing
 functions LR-21
 program structure SG-29
 program table, Multiple Terminal Manager M-21, M-31
 program termination, EXIO ID-157
 program/storage manager, Multiple Terminal Manager M-3
 program/task concepts SG-29
 PROGSTOP statement
 coding description LR-274
 overview LR-28, SG-31
 prompting and advance input CT-22
 protected field OR-322
 protocol, HCF transmission ID-184
 protocols,
 resynchronization SN-65
 PRT3780 utility SG-64
 overview SG-15, SG-64, SG-85
 PS command (\$IAMUT1) IG-169
 PSW (processor status word) MC-284, SG-275
 PSW bit settings MC-282
 PSW interrupt types MC-282
 punch data set formats SN-213
 punch data set, defining SN-213
 punch function, \$RJESNA SN-205
 PUNCHO command (\$RJESNA) SN-213
 PUNCHO/PUNCHS command (\$RJE2780/\$RJE3780) CT-340
 PUNCHS command (\$RJESNA) SN-213
 PUS, physical unit services
 control block SI-371, SI-372
 pushing a digital output address OR-368
 PUT command \$FONT OR-249
 full-screen text editor (\$FSEDIT) OR-259
 PUT Indexed Access Method CALL
 description CA-87, IG-127
 return codes CA-88, IG-128
 PUTCUR subroutine CA-480
 PUTDAT subroutine CA-481
 PUTDE Indexed Access Method CALL
 description CA-89, IG-130
 return codes CA-89, IG-131
 PUTEIT instruction
 coding description LR-276
 function SG-79
 overview LR-13
 return codes LR-278, MC-232

PUTFMT subroutine CA-485
PUTUP Indexed Access Method CALL
description CA-91, IG-133
return codes CA-92, IG-134
Px parameter naming operand LR-8,
SN-11

Q

QC command SN-137
QCB resource control block ID-45
QCB statement
coding description LR-282
overview LR-28
QD queue descriptor LR-24
QE queue entry overview LR-24
QEC command SN-137
QUALITY command \$DEBUG OR-160
QUESTION instruction
coding description LR-284
function SG-79, SG-80
overview LR-30
queuable resources SG-42
queue control block
See QCB statement
queue descriptor LR-24
queue entry
See QE queue entry
queue processing facilities SG-41
queue processing
instructions LR-24

R

RA command (\$TERMUT1) OR-465
random records IG-18
RATE parameter, SNAPU SN-105
RBN (relative block
number) CA-18, IG-40
RCB (Remote Management Utility
control block)
description ID-191
listing ID-366
RD command (\$INITDSK) OR-367
RDCURSOR instruction
coding description LR-287
function SG-79
overview LR-30
RDSCB parameter, NETINIT SN-20
RE command (\$COPY) OR-107
RE command (\$DISKUT1) OR-168
RE command (\$DIUTIL) CT-201
RE command (\$GPIBUT1) CT-139
RE command (\$HCFUT1) CT-370
RE command (\$IAMUT1) CA-113,
IG-170
RE command (\$SIS1UT1) CT-104
RE command (\$TERMUT1) OR-466
RE command (\$TERMUT2) OR-478
READ command (\$EDIT1/N) OR-208
READ IDCBC command LR-201
READ instruction
coding description LR-290
disk/diskette return
codes LR-293, MC-234
function SG-57
overview LR-16

tape post codes LR-293,
LR-296, MC-222
tape return codes LR-293,
LR-295, MC-235
read operations, HCF ID-184
READDATA command (\$HCFUT1) CT-368
READID IDCBC command LR-201
reading analog input OR-373
reading character image table from
4979 or 4978 OR-247
reading data set into work file
with \$EDIT1/N OR-207
with \$FSEDIT OR-268
reading digital input using
external sync OR-371
reading program
with \$UPDATE PP-108
with \$UPDATEH PP-119
reading records from host
(\$HCFUT1) CT-369
READOBJ command (\$HCFUT1) CT-369
READTEXT instruction
coding description LR-300
function SG-79
gather read operations CT-40
overview LR-30
reading unprotected
data CT-41, CT-43
return codes LR-253
using on 3101 terminals CT-43
ready task supervisor state ID-31
READ1 IDCBC command LR-201
READ80 command (\$HCFUT1) CT-369
realtime data member
\$PDS utility SG-325
RT \$DICOMP subcommand CT-183
receive buffers,
specifying SN-114
RECEIVE function, Remote Manage-
ment Utility
communications flow CT-271
internals ID-199
overview CT-269
sample program CT-302
receive operations (NETCTL) SN-51
receiving function management
headers SN-42
receiving messages SN-37
receiving messages, overview SN-8
receiving status SN-47, SN-51
RECLLEN parameter, NETGET SN-37
RECONNECT Multiple Terminal Manag-
er utility CA-306
record blocking, Remote Management
Utility CT-255
record exchange, Remote Management
Utility CT-253
record format, specifying in
DSD statement SM-74
record formats, explanation
fixed SM-77
fixed block SM-79
variable SM-81
variable blocked SM-81
record header, Remote Management
Utility CT-253
record locks, access
method AM-11, IM-12
directory IM-29
internals IM-25
record selection
advantages SM-6
examples SM-20
how to include SM-93

record size, determining SM-241
 record sizes, Host Communications Facility CT-349
 record summarization examples SM-29
 record/request processing AM-2, IM-2
 recording system release level LR-36, SG-239
 RECVBUF parameter, SNALU SN-114
 redirecting spool job(s) CA-415
 reformatting diskettes OR-119
 register conventions \$SIASM ID-282
 BSCAM processing ID-177
 common emulator setup routine ID-53
 EBCDIC to floating-point conversion ID-227
 terminal I/O support ID-99
 register, index LR-7
 register, software LR-7
 rejecting host initiated transactions SN-41
 REL command (\$HCFUT1) CT-370
 relating Series/1 support and applications SI-3
 relational statements LR-206
 relative block number (RBN) CA-18, IG-40
 definition of loading sequential file CA-37
 RELEASE Indexed Access Method CALL description CA-93, IG-136
 return codes CA-93, IG-136
 releasing space, \$PDS utility SG-335
 releasing spool job(s) CA-429
 releasing status record (\$HCFUT1) CT-370
 relocating program loader ID-15
 relocation dictionary, \$EDXASM ID-277
 RELQ command SN-137
 REMARK command (\$JOBUTIL) OR-394
 remote job entry utilities, BSC See \$RJE2780, \$RJE3780, \$PRT2780, and \$PRT3780
 remote job entry utility, SNA OR-420, SG-15, SN-201
 remote job entry, SNA SG-15, SG-85, SN-201
 Remote Management Utility CDRRM equates CT-330
 control block (RCB) description ID-191
 listing ID-366
 data areas ID-191
 defaults CT-320, ID-213
 error handling CT-327
 function table ID-191
 functions CT-250, ID-193, SG-64
 general equate listing ID-369
 installation CT-252
 logic flow ID-196
 modifying defaults CT-320
 module descriptions ID-217
 module list ID-212
 operation CT-257, ID-193
 overlay function process table ID-192
 overview CT-10, CT-249, SG-82
 problem determination ID-215
 program preparation ID-210
 requirements CT-250
 sample programs CT-300
 system generation considerations CT-251, SG-176
 tables ID-191
 removing breakpoints and trace ranges OR-156
 removing storage modules IG-206
 renaming member (\$DIUTIL) CT-201
 renaming volumes, RV \$INITDSK command OR-366
 RENUM command (\$FSEDIT) OR-278
 RENUM subcommand (\$EDIT1/N) OR-226
 reorganizing indexed file CA-61, IG-171
 report data member (\$PDS) SG-325
 reposition line pointer See moving line pointer
 REQDISCONT command SN-138
 request for data record, PASSTHRU function of \$RMU CT-293
 request record, Remote Management Utility CT-253
 request/response header SN-101
 request/response header, format SN-300
 requesting right-to-send SN-53
 reserved labels LR-5
 RESET command (\$RJESNA) SN-214
 RESET command, \$RJE2780/\$RJE3780 CT-341
 RESET control statement (\$EDXLINK) PP-56
 RESET instruction coding description LR-314
 overview LR-28, SG-31
 RESET option (\$EDXASM) PP-14
 RESET option (\$FSEDIT) OR-279
 reset punch data set name SN-214
 resetting device, Series/1 to Series/1 CT-104
 resetting GPIB adapter CT-140
 resetting indexed file parameters CA-113, IG-170
 resetting line commands, \$FSEDIT OR-279
 resident assembler routines, \$SIASM ID-281
 resident program loader ID-15
 resolution, graphics enhanced and standard ID-225
 resource control, supervisor ID-44
 restart, automatic SG-174
 restart, session SN-24
 restarting printer spooling CA-438
 restarting spool writer CA-431
 restore/dump volume, \$MOVEVOL OR-404
 restoring disk or disk volume from tape OR-452
 restoring 4974 to standard character set OR-478
 RESYNC parameter, NETINIT SN-21
 resynchronization data set SN-63
 resynchronization data set contents SN-64
 resynchronization data, disk SN-65

resynchronization data,
 storage SN-65
 resynchronization protocols SN-65
 resynchronization support, speci-
 fying SN-21
 retrieving Series/1 data
 set OR-208
 RETRY parameter, SNAPU SN-105
 return codes
 \$DISKUT3 MC-225, SG-343
 \$IMDATA subroutine CT-54,
 MC-239
 \$IMOPEN subroutine CT-47,
 MC-239
 \$IMPROT subroutine CT-52,
 MC-240
 \$PDS MC-226
 binary synchronous communi-
 cations CT-246, MC-227
 channel attach CT-440, MC-229
 checking SN-10
 CONVTB LR-70, MC-232
 CONVTD LR-75, MC-232
 description SG-232
 disk/diskette LR-295, MC-234
 EXIO instruction LR-141,
 MC-236
 EXIO interrupt LR-141, MC-237
 extended error MC-259, SN-87
 FADD LR-148, MC-238
 FDIVD LR-151, MC-238
 floating point MC-238
 FMULT LR-161, MC-238
 FSUB LR-179, MC-238
 general terminal I/O LR-254,
 MC-269
 GETEDIT LR-183, MC-232
 Host Communications
 Facility CT-373, MC-276
 Indexed Access Method
 DELETE CA-68, IG-102
 DISCONN CA-69, IG-104
 ENDSEQ CA-71, IG-106
 EXTRACT CA-74, IG-109
 GET CA-77, IG-113
 GETSEQ CA-80, IG-116
 LOAD CA-83, IG-120
 PROCESS CA-86, IG-125
 PUT CA-88, IG-128
 PUTDE CA-89, IG-131
 PUTUP CA-92, IG-134
 RELEASE CA-93, IG-136
 summary CA-96, IG-139,
 MC-241
 LOAD LR-221, MC-244
 Multiple Terminal
 Manager CA-278, MC-246
 NETBIND MC-260
 NETCLOSE MC-261
 NETCTL MC-251, SN-55
 NETGET MC-253, SN-44
 NETINIT MC-256, SN-28
 NETOPEN MC-262
 NETPUT MC-257, SN-36
 NETRECV MC-263
 NETSEND MC-264
 NETTERM MC-258, SN-62
 NETUBND MC-267
 printer spooling MC-249
 PRINTTEXT LR-253
 PUTEDIT LR-278, MC-232
 READ disk/diskette LR-290,
 MC-234
 READ tape LR-296, MC-235

 READTEXT LR-253
 SBI0 instruction MC-247,
 SG-367
 SDLC device MC-248, SN-113
 See also completion codes and
 post codes
 session termination MC-268
 Sort/Merge SM-135
 Systems Network
 Architecture MC-251
 tape LR-67
 TERMCTRL LR-253
 terminal I/O LR-306, MC-269
 terminal support
 FILEIO CA-247
 for all EDX terminals CT-28
 for DEVICE=ACCA CT-29,
 MC-270
 FTAB CA-255
 GETFMT CA-474
 OPENT CA-478
 PUTDAT CA-481
 PUTFMT CA-485
 SETFMT CA-270
 SETPAN CA-273
 summary CA-278, CA-488
 virtual terminals CT-206,
 MC-274
 WHEREAS LR-389
 WRITE disk/diskette LR-394,
 MC-234
 WRITE tape LR-394, MC-235
 return from immediate action
 routine (SUPEXIT) ID-40
 return from task level
 (SUPRTURN) ID-40
 RETURN instruction
 coding description LR-316
 overview LR-20, SG-41
 RETURN supervisor interface ID-51
 REW (rewind tape) LR-65, OR-449
 rewinding tape OR-449
 RH command (\$TERMUT1) OR-467
 RHDATA area SN-299
 RI command (\$BSCUT2) CT-243
 right-to-send, definition of SN-7
 right-to-send, granting SN-30
 right-to-send, requesting SN-53
 ring support, specifying SN-105
 RLOADER resident loader
 module description ID-72
 overlay programs ID-19
 overlay segments ID-20
 RO command (\$IAMUT1) CA-112,
 IG-171
 ROFF (rewind offline) LR-65
 roll screen, terminal support
 description CT-7, CT-15, CT-25
 differences from static
 screen CT-31
 sample roll screen
 program CT-25
 ROLLOFF command (\$COPYUT1) OR-113
 ROLLON command (\$COPYUT1) OR-112
 RP command (\$UPDATE) PP-108
 RP command (\$UPDATEH) PP-119
 RS command (\$GPIBUT1) CT-140
 RS command (\$SIS1UT1) CT-104
 RSHUTD command SN-138
 RSTATUS IDCB command LR-201
 RT command (\$TAPEUT1) OR-452
 RT function (\$PDS) SG-332
 RT subcommand (\$DICOMP) CT-183
 RTR command SN-138

RTYPE parameter, NETINIT SN-22
 RV command (\$INITDSK) OR-366
 RWI command (\$BSCUT2) CT-240
 RWIV command (\$BSCUT2) CT-245
 RWIVX command (\$BSCUT2) CT-244
 RWIX command (\$BSCUT2) CT-241
 RWIXMP command (\$BSCUT2) CT-242

S

SA subcommand (\$DICOMP) CT-183
 sample SNA application
 (SNAIMS) SN-263
 SAVE command (\$EDIT1/N) OR-227
 SAVE command (\$IMAGE) OR-337
 saving a control store OR-479
 saving a formatted screen OR-337
 saving current task status (TASK-
 SAVE) ID-44
 saving disk or diskette volume on
 tape OR-457
 saving graphics data CT-183
 saving image store OR-480
 saving session parameters SN-20
 saving storage and
 registers OR-491
 SB command (\$IOTEST) OR-372
 SBIO instruction
 coding description SG-367
 function SG-59, SG-360
 return codes MC-247, SG-369
 SBIOCB (sensor based I/O control
 block) ID-159
 SC command (\$TERMUT2) OR-479
 scatter write operation CT-37,
 CT-39, CT-42
 screen format builder utility
 See \$IMAGE utility
 screen formats, terminal support
 for 3101 CT-37
 for 4978/4979 CT-36
 sample for 3101 static screen
 program CT-79
 using \$IMOPEN to
 retrieve CT-41
 screen formatting/presentation
 support SN-171
 screen image manager, Multiple
 Terminal Manager M-7
 screen images
 buffer sizes CT-54
 retrieving and
 displaying CT-41
 storage format M-22
 using \$IMAGE subroutines CT-45
 SCREEN instruction
 coding description CT-157
 interprogram dialogue
 overview CT-157
 screen table, Multiple Terminal
 Manager M-21, M-31
 screens, terminal support
 logical screen
 See logical screen, terminal
 support
 roll screen
 See roll screen, terminal
 support
 static screen
 See static screen, terminal
 support

SCRNADD screen CA-313
 SCRNFIO1 screen CA-319
 SCRNFIO2 screen CA-320
 SCRNIPL screen CA-214
 SCRNMENU screen CA-216
 SCRNM1 screen CA-308
 SCRNM2 screen CA-309
 SCRNM3 screen CA-310
 SCRNM4 screen CA-311
 SCRNPMP screen CA-315
 SCRNPGM screen CA-314
 SCRNS volume CA-213
 SCRNSIGN screen CA-215
 SCRNSREP screen CA-317
 SCRNUPT screen CA-312
 scrolling, \$FSEDIT OR-261
 SCSS IDCB command LR-201
 SDLC buffers, specifying SN-112
 SDLC device address SN-105
 SDLC device I/O retries SN-107
 SDLC support options SN-108
 SDLC error logging SN-12
 SDLC log record, format SN-12
 SDLC return codes MC-248, SN-113
 SDLC time-outs, specifying
 clear-to-send SN-109
 data set ready SN-109
 data terminal ready SN-109
 default value SN-109
 hold line active SN-109
 SDT command SN-138
 SE command (\$HCFUT1) CT-370
 SE command (\$IAMUT1) CA-114,
 IG-173
 searching and dumping tape OR-437
 second-level index block
 (SIXB) CA-164, IG-35
 secondary control statement data
 set PP-48
 secondary index IG-67
 secondary index directory IG-68
 secondary index, Indexed Access
 Method
 action chart IM-27
 module descriptions IM-55
 secondary keys IG-67
 secondary logical unit type P
 formats and protocols
 advantages of Series/1 as SLU
 type P SN-169
 example of SLU type P/DPM
 configuration SN-169
 message
 resynchronization SN-168
 overview SN-165
 relation to MFS
 formatting SN-167
 selecting SLU type P
 support SN-166
 SLU type P without DPM SN-165
 secondary option menus, session
 manager OR-56
 sectioning of program
 modules LR-20
 sector size for diskette initial-
 ization \$DASDI utility OR-124
 disk (4963) OR-136
 sector, disk/diskette SG-46
 segmenting SN-131
 selecting paging IG-167
 selection, record
 example SM-20
 how to specify SM-95
 self-defining terms LR-5

send buffers, specifying SN-114
 send data, HX \$DICOMP
 subcommand CT-177
 SEND function, Remote Management
 Utility
 communications flow CT-275
 internals ID-198
 overview CT-273
 send/receive buffers, SDLC SN-106
 SENDBUF parameter, SNALU SN-114
 sending commands to host SN-211
 sending function management head-
 ers SN-32
 sending message to another termi-
 nal OR-487
 sending messages SN-29
 sending messages, overview SN-6
 sending partial messages SN-35
 sending record to host, Host Commu-
 nications Facility CT-364
 sense codes SN-305
 sensor based I/O
 assignments SG-359
 control block (SBIOCB) ID-159
 device data block (DDB)
 listing ID-358
 module descriptions ID-70
 statement overview SG-360
 storage requirements SG-220
 support summarized SG-20
 SENSORIO statement
 description SG-101
 function SG-59
 separator records SN-205
 sequence numbers SN-63
 sequencing instructions,
 program LR-21
 sequential access, Indexed Access
 Method IG-90
 serially reusable resource
 (SRR) ID-48, SG-42
 Series/1 layer elements SI-5
 Series/1 network activation proce-
 dure SN-309
 Series/1 screen
 formatting/presentation
 support SN-170
 Series/1 SNA base layer
 module cross reference SI-230
 overview SI-4
 program design language
 (PDL) SI-129
 visual table of contents
 (VTOC) SI-103
 Series/1 SNA Control Block Data
 Maps
 BIND parameter control block
 (BND) SI-282
 branch vector table
 (BVT) SI-286
 correlation table entry,
 CTE SI-290
 data control block, DCB SI-302
 data link buffer, DLB SI-317
 data link terminal, DLT SI-327
 data set description,
 DSD SI-334
 device descriptor block,
 DDB SI-306
 function address table,
 FAT SI-337
 immediate device control block,
 IDC SI-340
 LUE control block, LUE SI-342
 LUL control block, LUL SI-342,
 SI-352
 physical unit services control
 block, PUS SI-371, SI-372
 program list control block,
 PLC SI-369, SI-370
 SNA data area, DAT SI-294
 system services control block,
 LUS SI-366, SI-367
 Series/1 SNA instructions, over-
 view SN-6
 Series/1 SNA SDLC layer
 module cross reference SI-280
 overview SI-4
 program design language
 (PDL) SI-244
 visual table of contents
 (VTOC) SI-233
 Series/1 SNA subset SN-3
 Series/1 SNA system
 generation SN-103
 Series/1 SNA top layer
 module cross reference SI-100
 overview SI-4
 program design language
 (PDL) SI-24
 visual table of contents
 (VTOC) SI-11
 Series/1 to Series/1
 components CT-83
 data transfers CT-84
 error recovery CT-86
 exercise attachment
 See \$S1S1UT1 utility 86
 initiating processor CT-84
 internals
 interrupt processing ID-132
 request processing ID-129
 IPL function CT-86
 overview CT-7, CT-83, SG-19,
 SG-80
 posting an event control block
 (ECB) CT-87
 processor relationships
 (IPL) CT-84
 program synchronization CT-84
 programming
 considerations CT-84
 reconfiguring CT-85
 responding processor CT-84
 return codes MC-274
 sample programs CT-90
 storage requirements SG-219
 using direct I/O CT-85
 service request (SRQ) CT-116
 session activation SN-17
 session manager
 adding an option OR-80
 automatic creation/deletion of
 data sets
 \$SMDELET data set deletion
 control data set OR-51
 \$SMALLOC data set allo-
 cation control data
 set OR-50
 disk utilities OR-50
 building a \$JOBUTIL
 procedure OR-87
 communication utilities OR-66
 creating a new menu OR-81
 data management OR-59
 diagnostic utilities OR-68
 execute program
 utilities OR-64

graphics utility OR-63
 interactive interface
 (\$EDXLINK) PP-68
 invoking OR-46
 invoking a \$JOBUTIL
 procedure OR-87
 job stream processor
 utilities OR-64
 logon OR-47
 menus
 minimum partition size
 required OR-71
 operational overview OR-71
 option combinations OR-69
 overview OR-71
 primary option OR-55
 procedures OR-77
 program function keys OR-70
 program preparation
 utility OR-56
 programs OR-76
 secondary option OR-56
 session manager non-interactive
 interface (\$EDXLIN PP-68
 storage usage OR-71
 tailoring the session manager
 terminal utilities OR-61
 text editing utility OR-56
 updating a primary option
 menu OR-81
 updating a procedure OR-83
 utilities supported
 session parameters, saving SN-20
 session partner, definition
 of SN-6
 session restart SN-24, SN-143
 session termination return
 codes MC-268
 session termination,
 immediate SN-60
 session termination,
 orderly SN-60
 session termination,
 overview SN-10
 session, definition of SN-6
 session, PASSTHRU
 See PASSTHRU function
 session, termination SN-59
 sessions, sample \$RJESNA SN-218
 SESSPRM parameter, NETINIT SN-20
 SET command (\$HCFUT1) CT-370
 set time and date SG-61
 SETCUR, Multiple Terminal Manager
 CALL
 coding description CA-266
 internals M-9
 overview of CA-207
 SETEOD subroutine SG-354
 SETFMT, Multiple Terminal Manager
 CALL
 coding description CA-268
 internals M-10
 overview of CA-207
 return codes CA-270
 SETPAN, Multiple Terminal Manager
 CALL
 coding description CA-272
 internals M-9
 overview of CA-208
 return codes CA-273
 setting breakpoints and trace
 ranges OR-149
 setting date and time OR-36
 setting indexed file
 parameters CA-114, IG-173
 setting paging area size IG-168
 setting status (\$HCFUT1) CT-370
 setting tape offline OR-449
 setup procedure for
 \$JOBUTIL OR-376
 with \$JOBUTIL OR-378
 SG command (\$IOTEST) OR-372
 SHIFTL instruction
 coding description LR-318
 overview LR-14
 SHIFTR instruction
 coding description LR-320
 overview LR-14
 SHUTC command SN-138
 SHUTD command SN-139
 SHUTDOWN function, Remote Manage-
 ment Utility CT-276, ID-207
 SI command (\$TERMUT2) OR-480
 SIG command SN-139
 SIGNONFL file CA-293
 single task program SG-31
 SIXB (second-level index
 block) CA-164, IG-35
 SLE sublist element, \$EDXASM
 format ID-240
 in instruction parsing ID-242
 instruction description ID-255
 used in \$IDEF ID-267
 SLU bracket initiation SN-128
 SNA base modules SI-230
 SNA base support, overview SN-99
 SNA command and data flows
 controlling message
 flow SN-190
 establishing a session SN-176
 receiving messages SN-186
 sending messages SN-177
 terminating a session SN-194
 SNA control block
 relationship SI-281
 SNA data area, DAT SI-294
 SNA functions SI-2
 SNA return codes MC-250
 NETBIND MC-260
 return codes MC-251, MC-253,
 MC-256, MC-257, MC-258
 SNA SDLC layer modules SI-280
 SNA storage requirements SN-295
 SNA top layer modules SI-100
 SNADACT command SN-100
 SNAIMS, sample SNA
 application SN-263
 SNALU configuration statement
 parameter
 considerations SN-115
 parameter descriptions SN-114
 syntax SN-114
 SNAPU configuration statement
 parameter
 considerations SN-107
 parameter descriptions SN-104
 syntax SN-104
 software register LR-7
 software trace table SG-249
 Sort/Merge
 data interpretation SM-221
 debugging SM-179
 examples
 merge SM-16
 output record
 classification SM-6, SM-22

output record
 formatting SM-3, SM-18
 record selection SM-1, SM-5
 record summarization SM-22, SM-29
 simple sort SM-2, SM-14, SM-44
 improving performance SM-155
 initiation SM-131
 installation SM-295
 job preparation SM-39
 messages SM-247
 overview SG-13, SM-1
 sample jobs SM-159
 space requirements SM-239
 statements
 alternate collating sequence SM-87
 comments SM-68
 data set definition SM-71
 field SM-105
 header SM-60
 include/omit SM-93
 log SM-57
 user exit routines SM-137
 SORTA (record address sort)
 description SM-7
 dropping control fields SM-125
 functions supported SM-12
 input data sets
 number allowed SM-60
 restrictions SM-60
 inserting records
 input exit SM-137
 output exit SM-140
 output record size, determining SM-242
 specification statements allowed with SM-60
 specifying in header statement SM-60
 when to use SM-7
 sorting, how to write a sort job
 steps in preparing and running SM-44
 using \$JOBUTIL SM-42
 using \$L operator command SM-40
 using \$SMMAIN SM-48
 using LOAD (see also LOAD instruction) SM-44
 SORTR (record sort)
 drop/keep control fields SM-65
 example using SM-14
 functions supported SM-12
 output record size, determining SM-243
 rules for data fields SM-90
 specification statements allowed with SM-64
 specifying in header statement SM-60
 SORTRS (record summary sort)
 examples using SM-8
 functions supported SM-13
 keep/drop control fields SM-64
 output record size, determining SM-243
 rules with data fields SM-129
 specification statements allowed with SM-64
 specifying in header statement SM-60
 source program compiling SG-71
 source program entry and editing PP-1, SG-64
 source program line continuation using \$EDXASM LR-4
 SP command (\$CHANUT1) CT-410
 SP function (\$RJE2780/\$RJE3780) CT-341
 SPA (spool active control block)
 description ID-147
 listing ID-371
 space requirements, determining for output data set SM-241
 for work data set SM-239
 SPACE statement
 coding description LR-322
 overview LR-19
 SPC (spool terminal control block)
 description ID-148
 listing ID-374
 special PI bit/group OR-372
 specification data set SM-10
 specifications, data conversion LR-163
 conversion specifications
 See conversion LR-163
 specifications, in-storage SM-131
 specifying dynamic storage (\$EDX-LINK) PP-77
 SPECPI define special process interrupt SG-364
 SPECPIRT instruction
 coding description SG-376
 function SG-59, SG-360
 speed rate, modem SN-105
 SPJ (spool job control block)
 description ID-146
 listing ID-372
 splitting volume OR-355
 SPM (spool master control block)
 description ID-144
 listing ID-373
 spool active job control block (SPA)
 description ID-147
 listing ID-371
 SPOOL command (\$RJE2780/\$RJE3780) CT-341
 spool control record
 COBOL, use of CA-452
 format CA-450
 FORTRAN, use of CA-451
 functions CA-450
 PL/I use of CA-451
 spool data set
 allocating CA-455
 designating CA-437
 estimating size of CA-456
 spool device
 definition of CA-395
 designating CA-458
 spool device control block (SPW)
 description ID-149
 listing ID-371
 spool forms type code
 altering CA-415
 changing CA-407
 using CA-406
 spool I/O control block (IOSPTBL)
 description ID-151
 listing ID-371
 spool job
 definition of CA-396
 deleting CA-418
 displaying status CA-421

estimating print time CA-458
 holding CA-426
 keeping CA-427
 releasing CA-429
 status definition CA-396
 spool job control block (SPJ)
 description ID-146
 listing ID-372
 spool master control block (SPM)
 description ID-144
 listing ID-373
 spool printer assignment
 (\$RJESNA) SN-213
 spool return codes MC-249
 spool session
 definition of CA-396
 starting CA-401
 terminating CA-430
 spool terminal control block (SPC)
 description ID-148
 listing ID-374
 spool utility
 See \$SPLUT1 utility
 spool writer
 definition of CA-396
 displaying status CA-421
 functions CA-402
 procedure CA-402
 restarting CA-431
 starting CA-435
 stopping CA-433
 spool writer control block (WCB)
 description ID-150
 listing ID-374
 spooled output, \$RJESNA SN-204
 spooling
 See printer spooling
 SPW (spool device control block
 description ID-149
 listing ID-371
 SQ command (\$COPYUT1) OR-113
 SQRT instruction
 coding description LR-323
 overview LR-14
 SS command (\$DISKUT2) OR-190
 SSCPID parameter, NETHOST SN-18
 ST command (\$CHANUT1) CT-410
 ST command (\$DIUTIL) CT-201
 ST command (\$GPIBUT1) CT-140
 ST command (\$SIS1UT1) CT-105
 ST command (\$TAPEUT1) OR-457
 stand-alone dump diskette OR-120
 stand-alone dump,
 obtaining SG-240
 standard labels, tape
 EOF1 SG-295, SG-298
 EOV1 SG-295, SG-298
 fields SG-297
 HDR1 SG-295, SG-297
 header label SG-294
 layouts SG-296
 processing SG-296
 trailer label SG-294
 volume label SG-294
 VOL1 SG-295, SG-297
 start and termination procedure,
 \$DEBUG OR-145
 START, IDCB command LR-201
 START, PROGRAM statement
 operand LR-266
 starter system, supervisor
 description SG-167
 terminal initialization SG-213
 starting \$RJESNA SN-207
 starting spool writer CA-435
 starting spooling CA-401
 statement label LR-4
 statement locations, directory
 of LR-34
 static screen, terminal support
 description CT-7, CT-31
 designing for device independ-
 ence CT-38
 differences from roll
 screen CT-31
 3101 considerations CT-64
 3101 sample program CT-65
 4978/4979 sample program CT-60
 station ID, exchange SN-107
 status commands (\$HCFUT1) CT-370
 status data set, Host Communi-
 cations Facility CT-350
 status display, WHERE
 \$DEBUG OR-162
 status obtaining error
 with \$GPIBUT1 CT-140
 with \$SIS1UT1 CT-105
 status record, Remote Management
 Utility CT-253, CT-299
 STATUS statement
 coding description LR-325
 overview LR-13
 status, receiving SN-47, SN-51
 STAXID parameter, SNAPU SN-107
 STIMER instruction
 coding description LR-328
 in Series/1 to Series/1 error
 recovery CT-87
 overview LR-32, SG-42
 with PASSTHRU function CT-295
 stopping spool writer CA-433
 stopping spooling facility CA-430
 storage environment, access
 method AM-1, IM-1
 storage environment, Multiple
 Terminal Manager M-11
 storage estimating
 application program
 size SG-225
 supervisor size SG-217
 utility program size SG-223
 storage management
 address relocation
 translator SG-43
 allocating ID-24
 description SG-43
 design feature SG-16
 storage map generation SG-206
 storage map, Indexed Access
 Method CA-15, IG-204
 storage map, resident
 loader ID-26
 storage maps (\$S1ASM) ID-286
 storage requirements, SNA SN-295
 storage usage during IPL ID-6
 storage usage during program
 load ID-17
 storage-resident loader
 (RLOADER) ID-15
 store record, \$PDS utility SG-335
 strings, relational
 statement LR-206
 STSN command SN-65, SN-139
 STSN processing
 PLU to SLU flow (with
 resync) SN-76
 return codes to NETINIT SN-85
 sessions without resync SN-84

SLU to PLU flow (with resync) SN-68
 STSN processing, Series/1 SN-67
 SU command (\$GPIBUT1) CT-141
 SU command (\$HCFUT1) CT-371
 SUBMIT command (\$EDIT1/N) OR-209
 SUBMIT command (\$FSEDIT) OR-270
 SUBMIT command (\$RJESNA) SN-214
 SUBMIT command (\$RJE2780/\$RJE3780) CT-342
 submitting job streams SN-214
 submitting job to host (\$HCFUT1) CT-371
 SUBMITX command (\$RJESNA) SN-214
 SUBMITX command (\$RJE2780/\$RJE3780) CT-342
 SUBROUT statement
 coding description LR-330
 in I/O error exit routine SM-143
 in input exit routine SM-137
 in output exit routine SM-141
 overview LR-20, SG-41
 subroutines
 \$DISKUT3 SG-339
 \$IMAGE CT-44
 \$PDS SG-321
 DSOPEN SG-347
 overview SG-40
 SETEOD SG-354
 subset, Series/1 SNA SN-3
 SUBTRACT instruction
 coding description LR-332
 overview LR-14
 precision table LR-333
 summary data field
 example SM-29
 extending SM-32
 how to specify SM-126
 length, minimum/maximum SM-112
 maximum number SM-125
 rules for specifying SM-125
 summary tagalong sort
 See SORTRS (record summary sort)
 supervisor/emulator
 calling control functions ID-49
 class interrupt vector table ID-11
 communications vector table description ID-12
 listings ID-319
 concepts SG-29
 control block pointers ID-12
 design features SG-15, SG-16
 device vector table ID-12
 emulator command table description ID-14
 listing ID-323
 entry points SG-465
 entry routines ID-37
 exit routines ID-40
 features SG-15, SG-16
 fixed storage area ID-11
 functions ID-34
 generation SG-167
 interface routines GPIB ID-119
 printer spooling ID-139
 Series/1 to Series/1 ID-129
 task supervisor ID-50
 maintaining more than one SG-193

module descriptions ID-57
 module names SG-465
 module overview ID-58
 overview SG-5, SG-29
 PASSTHRU session with CT-281
 referencing storage locations ID-13
 service routines ID-44
 size, estimating SG-217
 storage requirements SG-218
 task supervisor description ID-29
 work area ID-14
 SUPEXIT routine ID-40
 support for optional features LR-13
 SUPRTURN routine ID-40
 surface analysing tape OR-439
 suspending \$GPIBUT1 CT-141
 suspending message flow SN-53
 suspending transmission SN-150
 SV command (\$INITDSK) OR-355
 SVC routine ID-38
 SVCABEND routine ID-40
 SVCBUF supervisor buffer ID-38
 SVCI routine ID-38
 switched major node definition operands SN-324
 sample SN-325
 symbol dictionary, \$EDXASM ID-277
 symbolic address (disk,tape) LR-10
 symbolic disk/tape I/O assignments LR-10
 symbolic diskette LR-10
 symbolic reference to terminals SG-116
 symbolic terminal I/O LR-10
 symbols (EXTRN) LR-145
 symbols (WXTRN) LR-399
 syntax rules LR-4
 SYSGEN
 See system generation
 system alternate logging device SG-116
 system AUTOCALL data set (\$EDXLINK) PP-75
 system class interrupt vector table ID-11
 system common area ID-13
 system common data area (\$SYSCOM) SG-11, SG-172
 system communications vector table description ID-12
 listing ID-319
 system configuration for program preparation SG-24
 minimum for execution SG-23
 minimum licensed programs SG-24
 statements SG-89
 system control blocks SG-44
 system data tables, EDXSYS ID-67
 system device vector table ID-12
 system emulator command table description ID-14
 listing ID-323
 system generation
 customizing the starter system SG-194
 procedure SG-179
 system installation SG-27
 system logging device SG-116
 system printer SG-116

system products SG-3
 system program check MC-282,
 SG-275
 system release level,
 recording LR-36, SG-239
 system reserved labels LR-5
 system services control block,
 LUS SI-366, SI-367
 SYSTEM statement SG-103
 system task supervisor work
 area ID-14
 System/370 Channel Attach
 See Channel Attach
 System/370 Program Preparation
 Facility SG-7

T

TA command (\$TAPEUT1) OR-461
 tables, Multiple Terminal Manager
 file table M-19, M-35
 program table M-21, M-31
 screen table M-21, M-31
 terminal table M-19, M-31
 tables, parameter equate LR-431
 tabs setting
 horizontal (\$IMAGE) OR-331
 vertical (\$IMAGE) OR-339
 with \$EDIT1/N OR-228
 TABSET command (\$EDIT1/N) OR-228
 tagalong sort
 See SORTR (record sort)
 tailoring the Indexed Access Meth-
 od IG-144
 tailoring your supervisor SG-194
 tape
 bypass label processing SG-300
 characteristics and
 layout SG-55
 CONTROL LR-64
 data set LR-26
 definitions for data
 sets LR-26
 end-of-tape (EOT) LR-26
 examples SG-310
 I/O instructions LR-26
 internals ID-91
 load point (BOT) LR-26
 post codes MC-222
 READ LR-290
 record LR-26
 return codes LR-67, MC-235
 statement SG-111
 symbolic addressing LR-10
 used in Version 3
 conversion SG-439
 volume LR-26
 WRITE LR-391
 tape control commands OR-449
 tape data set
 end-of-data requirements SM-82
 rules for specifying SM-72
 tape data set control block ID-93
 tape device data block (TDB)
 description ID-92
 listing ID-377
 tape labels
 external SG-293
 internal SG-294
 tape standard label
 fields SG-297

layout SG-296
 processing SG-296
 support SG-294
 TAPE statement
 coding SG-111
 description SG-111
 function SG-58
 tape support SG-17
 tape utility, \$TAPEUT1 OR-431
 tape volumes, defining SG-304
 tape, non-labeled
 layout SG-299
 processing SG-300
 support SG-299
 tapemark LR-64
 task
 active/ready table ID-41
 check and trap handling SG-231
 control ID-44
 definition and control func-
 tions SG-30
 dispatching ID-42
 error exit facility SG-41,
 SG-253
 error exit routine LR-270
 error handling LR-31
 execution ID-30
 execution states ID-30, SG-37
 facilities SG-30
 linkage conventions SG-256
 multiple-task program SG-33
 overview LR-28, SG-30
 single-task program SG-31
 states SG-37
 structure SG-30
 switching ID-42, SG-30
 synchronization and
 control ID-44, SG-30
 task code words LR-9
 task control block
 See TCB (task control block)
 task error exit SN-11, SN-217
 TASK statement
 coding description LR-335
 overview LR-28, SG-31
 priority LR-335
 task supervisor
 functions ID-34
 internals ID-29
 work area ID-14
 TASKSAVE supervisor service
 routine ID-44
 TCB (task control block)
 description ID-32
 listing ID-375
 with QCB ID-46
 TCBEQU LR-431
 TCBGET instruction
 coding description LR-338
 overview LR-28
 TCBPUT instruction
 coding description LR-340
 overview LR-28
 TD command (\$PDS) SG-332
 TD subcommand (\$DICOMP) CT-184
 TDB (tape device data block)
 description ID-92
 EDXTIO/EDXTIOU module ID-68
 listing ID-377
 virtual I/O ID-111
 TEB (terminal environment block)
 description CA-220, M-17
 equates M-32
 Tektronix

devices supported SG-17
 support for digital I/O CT-465
 teletypewriter adapter SG-148
 TERM-SELF command SN-139
 TERMCTRL instruction
 ACCA attached devices LR-361
 coding descriptions LR-342
 function to terminal comparison
 table LR-343
 General Purpose Interface
 Bus LR-371
 overview LR-30
 return codes LR-306
 Series/1 To Series/1 LR-367
 teletypewriter attached
 devices LR-364
 use on 3101 terminals CT-44
 virtual terminal LR-365
 2741 Communications
 Terminal LR-344
 3101 Display Terminal (block
 mode) LR-345
 4013 Graphics Terminal LR-349
 4973 Printer LR-351
 4974 Printer LR-352
 4975 Printer LR-354
 4978 Display LR-355
 4979 Display LR-359
 TERMERR operand
 of PROGRAM LR-270
 of TASK LR-335
 with end-of-forms LR-31
 terminal
 #7850 teletypewriter
 adapter SG-148
 ACCA support LR-361
 adding support for new I/O
 terminals ID-113
 characteristics CT-3
 data representation LR-31
 defining and accessing CT-16
 definition CT-3, SG-79
 device independence
 between display stations and
 printers CT-7
 between 4978/4979 and
 3101 CT-7
 for static screens CT-38,
 CT-39
 electrical interfaces CT-5
 error handling LR-30
 input LR-32
 list of EDX terminals CT-4
 output LR-31
 relationship to a
 program CT-16
 return codes LR-253, MC-269
 sample I/O error
 program CA-352
 sample signon program CA-352
 specialized capabilities
 general purpose interface
 bus CT-7
 graphics CT-8
 Series/1 to Series/1 attach-
 ment CT-7
 virtual terminals CT-8
 support summarized SG-17
 terminal configuration utility,
 \$TERMUT1 OR-462
 terminal control block
 See CCB (terminal control
 block)
 terminal environment block

See TEB (terminal environment
 block)
 TERMINAL file CA-212
 terminal I/O
 high-level language subroutines
 COBOL sample program CA-492
 FORTRAN sample
 program CA-499
 introduction CA-464
 overview CA-465
 performance CA-506
 PL/I sample program CA-495
 preparing programs CA-503
 internal design ID-98
 problem determination ID-112
 return codes MC-269
 See also terminal support
 support routines ID-106
 virtual ID-111
 terminal message sending utility,
 \$TERMUT3 OR-487
 terminal page formatting param-
 eters OR-463
 terminal response mode SN-160
 terminal servers, Multiple Termi-
 nal Manager M-6
 terminal session manager
 See session manager
 TERMINAL statement
 coding by device
 ACCA CT-80, SG-136
 GPIB SG-160
 PROC CT-81, SG-154
 Series/1 to Series/1 SG-163
 TTY CT-80, SG-148
 virtual terminal SG-158
 2741 SG-120
 4013 SG-123
 4973/4974 SG-127
 4975 SG-129
 4978/4979 SG-132
 description SG-113
 DEVICE parameter CT-4
 device-dependent
 operands SG-116
 examples CT-16
 function SG-79
 terminal control
 function SG-80
 use in defining virtual termi-
 nals CT-204
 use of CT-3, CT-4
 terminal states, Multiple Terminal
 Manager M-39
 terminal support
 definition CT-3
 EDL instructions
 for controlling a
 terminal CT-18
 for defining and accessing a
 terminal CT-18
 for formatting data CT-19
 output considerations CT-23
 overview CT-3, SG-77
 sample roll screen
 program CT-25
 3101 sample static screen
 program CT-65
 4978/4979 sample static screen
 program CT-60
 terminal table, Multiple Terminal
 Manager M-19, M-31
 terminal utilities (session manag-
 er) OR-61

terminals connected via digital I/O CT-467
 terminating \$RJESNA SN-211, SN-216
 terminating a session SN-59
 terminating GPIB operation CT-134
 terminating Remote Management Utility CT-276
 terminating transmission, \$RJESNA SN-210
 test display (\$DICOMP) CT-168
 testing BSC conversational transparent mode CT-244
 testing generated report or graphics profile member CT-168
 testing label types OR-439
 testing process interrupt OR-372
 text editing utilities
 edit dataset subroutine examples ID-388
 full screen-editor \$FSEDIT OR-259
 line editors, \$EDIT1/N OR-199
 overview SG-64
 work area internals ID-383
 text record data area ID-192
 text record, PASSTHRU function of \$RMU CT-289
 TEXT statement
 coding description LR-376
 overview LR-13
 text wrapping, WRAP function of \$RMU CT-279
 THRESH parameter, SNAPU SN-106
 threshold considerations SN-110
 threshold value SN-109
 time of day, GETTIME instruction LR-188
 time/date
 display, \$W operator command OR-44
 set, \$T operator command OR-36
 set, automatic initialization facility SG-175
 timer control LR-32
 TIMER statement SG-42, SG-165
 timer storage requirements SG-219
 timer support SG-20
 timing instructions LR-32, SG-42
 TITLE statement
 coding description LR-380
 overview LR-19
 TOCTS parameter, SNAPU SN-106
 TODSR parameter, SNAPU SN-105
 TODTR parameter, SNAPU SN-105
 TOHLA parameter, SNAPU SN-106
 TOP subcommand (\$EDIT1/N) OR-232
 TP instructions (Host Communications Facility)
 CLOSE CT-355
 coding description CT-353
 FETCH CT-356
 OPENIN CT-357
 OPENOUT CT-358
 READ CT-359
 RELEASE CT-360
 SET CT-361
 subcommand internals ID-182
 SUBMIT CT-362
 TIMEDATE CT-363
 use in examples CT-365
 WRITE CT-364
 TR command (\$CHANUT1) CT-410

trace printing utility for BSC CT-237
 trace ranges and breakpoint settings OR-149
 trace table, software SG-249
 trace utility for BSC CT-235
 transaction retry, requesting SN-41
 transactions, definition of SN-6
 transfer rates for data, Host Communications Facility CT-353
 transferring an object module SG-72
 transferring data set from host (\$HCFUT1) CT-368
 transferring data set to host SEND function, Remote Management Utility CT-273
 WR command (\$HCFUT1) CT-372
 transient program loader ID-15
 transmission codes SG-114
 transmission header, format SN-301
 transmission protocol, HCF ID-184
 transmission, suspending SN-150
 TRAPDUMP attention command (\$TRAP) OR-491
 TRAPEND attention command (\$TRAP) OR-491
 TRAPOFF attention command (\$TRAP) OR-491
 TRAPON attention command (\$TRAP) OR-491
 TS profiles 3 and 4 supported commands SN-119
 TYPE parameter, NETCTL SN-49
 TYPE= return codes (NETCTL) SN-57
 types of session termination SN-60

U

UE command (\$IAMUT1) IG-158
 UN command (\$IAMUT1) CA-120, IG-184
 UNBIND command SN-139
 unblocked sequential IG-164
 unconditional force data types allowed SM-109
 definition SM-117
 example SM-22
 how to specify SM-117
 keep/drop from output record SM-64
 number, maximum SM-121
 required entries in field statement SM-109
 rules for SM-116
 unconditional session termination SN-60
 undefined length records, tape SG-303
 units (basic) of information SN-297
 unloading indexed file CA-120, IG-184
 unprotected field OR-322
 UP subcommand (\$EDIT1/N) OR-233
 update indicator, automatic IG-69
 update utility

\$UPDATE convert object
 program PP-105
 \$UPDATEH convert host object
 program PP-117
 updating a session manager
 menu OR-81
 updating directory IG-158
 user defined data member, \$PDS
 utility SG-326
 user exit routines SM-33
 keep/drop from output record
 specifying in header SM-64
 selecting the invocation facility
 SM-39
 user initialization modules ID-10
 USER instruction
 coding description LR-381
 effect on ENDPROG LR-116
 overview LR-20
 using SELB to connect
 support SI-3
 utilities
 listed by type SG-62
 overview SG-6
 utility program size SG-223

V

VA command (\$DICOMP) CT-183
 VA function (\$PDS) SG-328
 valid arithmetic operators LR-4
 variable blocked record
 format SM-81
 variable length records,
 tape SG-303
 variable names LR-4
 variable record format SM-81
 variable unblocked record
 format SM-80
 varying device offline OR-38
 varying device online OR-40
 VDE (volume descriptor
 entry ID-85
 vector addition (ADDV) LR-41
 vector data manipulation LR-15
 verification, requesting
 message SN-30
 VERIFY parameter, NETPUT SN-30
 VERIFY subcommand
 (\$EDIT1/N) OR-234
 verifying disk or diskette
 dataset OR-360
 verifying identification of
 systems, \$RMU CT-267
 verifying tape executing
 correctly OR-439
 tape, \$TAPEUT1 utility OR-442
 vertical tabs, defining OR-339
 virtual terminals
 accessing the virtual
 terminal CT-205
 defined by TERMINAL
 statement SG-158
 description CT-203
 establishing the
 connection CT-204
 examples of use CT-203
 internals ID-111
 interprogram dialogue CT-206
 loading from a virtual
 terminal CT-204

return codes CT-206, MC-274
 sample programs CT-207
 storage requirements SG-219
 TERMINAL statement
 example SG-159
 using TERMINAL to
 define CT-204
 volume
 accessing ID-82
 allocating OR-353
 definitions
 (disk/diskette) LR-17
 deleting OR-365
 directory entry
 listings ID-347
 directory internals ID-77
 dump/restore, \$MOVEVOL OR-404
 initializing OR-357
 listing OR-362
 renaming OR-366
 split OR-355
 verifying OR-360
 VOLUME control statement
 (\$EDXLINK) PP-56
 volume descriptor entry
 (VDE) ID-85
 volume labels
 logical SG-47
 physical SG-46
 VOL1 tape label SG-295
 VTAB command (\$INITDSK) OR-339
 VTAM considerations
 block number SN-293
 network deactivation SN-293
 pacing SN-293
 station address SN-293
 VTOCs for SDLC layer modules
 CDCABNO SI-238
 CDCADDT0 SI-238
 CDCATTN SI-238
 CDCBLDDC SI-238
 CDCCLS00 SI-239
 CDCCMDPR SI-239
 CDCCMDR SI-239
 CDCCNTBU SI-240
 CDCCNTRL SI-240
 CDCDISAB SI-240
 CDCENABL SI-240
 CDCIHT SI-241
 CDCISXMT SI-242
 CDCNOI SI-243
 CDCOIO SI-243
 CDCPERME SI-243
 CDCPURGE SI-243
 CDCRCSS SI-243
 CDCSETRE SI-243
 SDLC abnormal/error
 handling SI-235
 SDLC close SI-237
 SDLC interrupt handler SI-234
 SDLC issue command to
 adapter SI-236
 SDLC read cycle steal status
 processing SI-236
 SDLC receive SI-235
 SDLC send SI-235
 SDLC stop device SI-237
 VTOCs for Series/1 SNA base layer
 modules
 \$NETCMD, interface to SNA
 modules SI-107
 \$SNA, SNA load program SI-128
 CDDACTLU, ACT LU command
 processor SI-108

CDDBETBR, between brackets
 check SI-108
 CDDBINDO, bind command process-
 or SI-108
 CDDBNDC, process basic bind
 checking SI-108
 CDDBRSP, BIND response proces-
 sor SI-109
 CDDBSMGR, bracket state
 manager SI-109
 CDDCLEAR, reset data traffic
 subtree SI-109
 CDDCLS00, SNA close SI-110
 CDDCOMSC, common session
 control processor SI-111
 CDDCPMGS, CP manager send for
 LU-LU sessions SI-111
 CDDCPMGX, process expedited
 flow queue SI-111
 CDDCP2S, pacing queue
 processor SI-111
 CDDCTDLT, correlation table
 entry delete processor SI-112
 CDDCTGEN, RSP correlation table
 entry generation SI-112
 CDDCTPRC, correlation table
 entry processor SI-112
 CDDDRQRV, DFC commands process-
 or SI-112
 CDDEXRCV, convert RQ to
 EXR SI-113
 CDDFISC, FI for session control
 processor SI-113
 CDDGTBFS, allocate buffer
 pool SI-113
 CDDHXMGR, half-duplex
 manager SI-114
 CDDIMMSG, inbound message
 processor SI-114
 CDDISEG, queues and releases
 segmented buffers SI-115
 CDDLUPRG, purge queues SI-115
 CDDMRQCV, create negative RSP
 and SEND SI-115
 CDDMRSCV, change RCV/SEND buff-
 er count SI-115
 CDDNTBND, send user bind
 response SI-116
 CDDNTDAC, network
 deactivation SI-116
 CDDNTRCV, receive
 message SI-116
 CDDNTSND, send message SI-117
 CDDOPN00, SNA NETOPEN phase
 0 SI-118
 CDDOPN01, SNA NETOPEN phase
 1 SI-118
 CDDOPN02, SNA NETOPEN phase
 2 SI-118
 CDDOPN03, SNA NETOPEN phase
 3 SI-118
 CDDOPN04, SNA NETOPEN phase
 4 SI-118
 CDDPBUFF, post send
 event SI-119
 CDDPCLUR, SSCP-LU command
 processor SI-119
 CDDPCNGS, path control no-good
 message processor SI-119
 CDDPCPUR, PU message
 processor SI-119
 CDDPC20T, outbound path
 control SI-120

CDDPSV00, perform physical unit
 services SI-120
 CDDPUSIN, process write buffers
 and SDLC commands SI-120
 CDDQCHK, check all network
 queues SI-120
 CDDQOFF, remove from FIFO
 queue SI-121
 CDDQON, insert on FIFO
 queue SI-121
 CDDRECVY, half-duplex
 recovery SI-121
 CDDRELBL, release link
 buffer SI-122
 CDDRELBS, release
 buffer SI-122
 CDDRRQRCV, request received
 processor SI-123
 CDDRRRSZ, convert RQ to
 EXR SI-124
 CDDRRSNN, normal request send
 processor SI-124
 CDDRRSNX, expedited request
 send processor SI-124
 CDDRSRCV, response received
 processor SI-125
 CDDRSSNN, normal response send
 processor SI-126
 CDDRSSNX, expedited outbound
 response processor SI-126
 CDDSELB, interface to
 \$NETCMD SI-126
 CDDSESL, session
 cleanup SI-126
 CDDSTPLK, halt SDLC SI-127
 CDDSVCCR, process
 NETSEND/NETRCV
 requests SI-127
 CDDTCB, find active TCB SI-127
 CDDUBND, send term-self SI-127
 CDDWTUBD, wait for
 UNBIND SI-128
 CDDZPNMS, convert to minus
 response SI-128
 NETBIND SI-105
 NETCLOSE SI-107
 NETOPEN SI-104
 NETRCV SI-106
 VTOCs for Series/1 SNA Top Layer
 Modules
 CDDFANE SI-13
 CDEATTN SI-19
 CDECD SI-20
 CDECHKRU SI-19
 CDECHK00 SI-20
 CDECLEAN SI-20
 CDEDEQ SI-21
 CDEDEQAB SI-21
 CDEENQ SI-21
 CDEEXIT SI-21
 CDEFBUFF SI-21
 CDEGBUFF SI-22
 CDEINIT1 - NETINIT phase
 1 SI-15
 CDEPOST SI-22
 CDEPUT0 - NETPUT phase
 0 SI-17
 CDEPUT1 - NETPUT phase 1 SI-17
 CDERCVRU SI-22
 CDEREJ SI-22
 CDERSP SI-23
 CDERSYNC SI-23
 CDESNDRU SI-23

CDESTSN - NETINIT phase
 1 SI-16
 CDETERM0 - NETTERM phase
 1 SI-18
 CDETERM1 - NETTERM phase
 1 SI-18
 NETCTL phase 0 SI-12
 NETCTL phase 1 SI-12
 NETGET phase 0 SI-14
 NETGET phase 1 SI-14
 NETINIT phase 0 SI-15
 VV command (\$INITDSK) OR-360

W

wait state, put program in, WS
 \$IOTEST command OR-368
 WAIT statement
 coding description LR-385
 overview LR-28, SG-31
 use of WAIT KEY in terminal
 support CT-21
 WAIT supervisor function ID-36
 waiting, task supervisor
 state ID-31
 WAITKY subroutine CA-487
 warning message option SM-66
 WCB (spool writer control block)
 description ID-150
 listing ID-374
 WE command (\$COPY) OR-109
 WHERE command (\$DEBUG) OR-162
 WHERE instruction
 coding description LR-388
 overview LR-28
 return codes LR-389
 WI command (\$BSCUT2) CT-243
 WIX command (\$BSCUT2) CT-243
 word boundary requirement
 DO LR-21
 IF LR-21
 PROGRAM LR-266
 work areas, Multiple Terminal
 Manager
 buffer areas M-23, M-37
 common area M-16, M-34
 MTMSTORE work file M-20
 work data set SM-9
 \$EDXASM PP-11
 \$EDXASM internals ID-275
 \$EDXLINK PP-48
 \$LINK PP-86
 \$\$IASM PP-28
 \$\$IASM internals ID-284
 work file (MTMSTORE), Multiple
 Terminal Manager M-20
 workstation features SN-203
 workstation function priorities,
 \$RJESNA SN-205
 workstation functions
 card reader SN-204
 console SN-203
 printer SN-204
 punch SN-205
 workstation, defining
 Series/1 SN-225
 WR command (\$GPIBUT1) CT-142
 WR command (\$HCFUT1) CT-372

WR command (\$SISUT1) CT-105
 WRAP function, Remote Management
 Utility CT-279, ID-202
 WRITE command (\$FSEDIT1/N) OR-210
 WRITE command \$EDIT1 OR-210
 WRITE \$EDIT1 command
 WRITE OR-210
 WRITE disk/diskette I/O instruc-
 tion
 coding description LR-391
 disk/diskette return
 codes LR-396, MC-234
 function SG-57, SG-58
 IDCB command LR-201
 overview LR-16
 return codes SN-96
 tape post codes LR-397
 tape return codes LR-397,
 MC-234
 WRITE IDCB command LR-201
 WRITE instruction
 WRITE Multiple Terminal Manager
 CALL
 coding description CA-275
 overview of CA-208
 write operations, HCF ID-184
 WRITE option (\$FSEDIT) OR-269
 WRITE, Multiple Terminal Manager
 CALL
 internals M-10
 writer control block, spool (WCB)
 description ID-149
 listing ID-374
 WRITE1 IDCB command LR-201
 writing data to the GPIB
 adapter CT-142
 writing data, Series/1 to
 Series/1 CT-105
 writing decompression
 routines SN-223
 writing IPL text OR-358
 WS command (\$IOTEST) OR-368
 WTM (write tape mark) LR-65
 WXTRN statement
 coding description LR-399
 overview LR-20

X

X.21 protocol CT-11
 X-type format LR-170
 XI command (\$IOTEST) OR-371
 XO command (\$IOTEST) OR-370
 XYPLOT instruction
 coding description CT-158
 overview CT-158

Y

YTPLOT instruction
 coding description CT-159
 overview CT-159

1

1024-byte sectors
 alternate load SG-209
 I/O module description ID-70
 storage requirements SG-218
 taking advantage of SG-208
1310 Multifunction Attachment
 defined by ADAPTER
 statement SG-90
 description CT-9, CT-467
 use in BSC CT-11
 with 3101 Display
 Terminal CT-477
1610 asynchronous communications
 single line controller
 attachment with 3101 Display
 Terminal CT-471
 considerations for attachment
 of devices SG-113
 considerations when
 DEVICE=PROC CT-81

2

2091 asynchronous communications
 eight line controller
 attachment with 3101 Display
 Terminal CT-471
 considerations for attachment
 of devices SG-113
2092 asynchronous communications
 four line adapter
 attachment with 3101 Display
 Terminal CT-471
 considerations for attachment
 of devices SG-113
2095 feature programmable eight
 line controller
 attachment with 3101 Display
 Terminal CT-471
 considerations for attachment
 of devices SG-113
2096 feature programmable four
 line adapter
 attachment with 3101 Display
 Terminal CT-471
 considerations for attachment
 of devices SG-113
2741 Communications Terminal
 as an EDX terminal CT-5
 defined by TERMINAL
 statement SG-120
 storage requirements SG-219
 TERMINAL statement
 example SG-122

3

3101 Display Terminal
 as an EDX terminal CT-5, CT-6
 attribute characters CT-36,
 OR-326
 block mode CT-474, OR-259
 character mode CT-470
 data stream CT-36

defined by TERMINAL
 statement SG-136, SG-148
 interface with Multiple Termi-
 nal Manager CA-202
 PF key support CT-20
 reading modified data CT-57
 SEND key CT-478
 setup switch settings
 block mode CT-474
 character mode CT-470
 static screen
 considerations CT-64
 static screen sample
 program CT-65
 storage requirements SG-219
 TERMINAL statement
 examples SG-144, SG-152
 transmitting data from CT-36
3585 4979 Display Station attach-
 ment SG-132

4

4952 Processor
 maximum partitions SG-43
 timer feature SG-42
4953 Processor
 maximum partitions SG-43
 timer feature SG-42
4955 Processor
 maximum partitions SG-43
 timer feature SG-42
4962 Disk Storage Unit storage
 capacity SG-285
4963 Disk Storage Unit storage
 capacity SG-285
4964 Diskette Storage Unit
 description SG-51
 part of minimum
 configuration SG-23
 required for program prepara-
 tion SG-24
 storage requirements SG-218
4965 Diskette Storage Unit
 description SG-51
 part of minimum
 configuration SG-23
 required for program prepara-
 tion SG-24
 storage requirements SG-218
4966 Diskette Magazine Unit
 part of minimum
 configuration SG-23
 required for program prepara-
 tion SG-24
 storage requirements SG-218
4969 Magnetic Tape
 Subsystem SG-293
4973 Line Printer
 as an EDX terminal CT-5
 defined by TERMINAL
 statement SG-127
 end of forms CT-24
 storage requirements SG-218
 TERMINAL statement
 example SG-128
4974 Matrix Printer
 as an EDX terminal CT-5
 defined by TERMINAL
 statement SG-127
 end of forms CT-24

modify image/control store,
 \$FONT utility OR-242
 restore to character set, RE
 \$TERMUT2 OR-478
 storage requirements SG-218
 TERMINAL statement
 example SG-128
 4975 Printer
 as an EDX terminal CT-6
 defined by TERMINAL
 statement SG-129
 end of forms CT-24
 storage requirements SG-218
 TERMINAL statement
 example SG-131
 4978 Display Station
 \$FONT utility OR-468
 as an EDX terminal CT-5
 control/image store SG-212
 defined by TERMINAL
 statement SG-132
 modify image/control store
 static screen sample
 program CT-60
 storage requirements SG-218
 TERMINAL statement
 example SG-135
 4979 Display Station
 as an EDX terminal CT-5
 defined by TERMINAL
 statement SG-132
 static screen sample
 program CT-60
 storage requirements SG-218
 TERMINAL statement
 example SG-135
 4993 channel attach device CT-380

5

5230 Data Collection
 Interactive SG-13
 5620 4974 matrix printer attach-
 ment
 defined by TERMINAL
 statement SG-127

5630 4973 line printer attachment
 defined by TERMINAL
 statement SG-127
 5719-AM3
 See Indexed Access Method
 5719-AM4
 See Indexed Access Method
 5719-ASA
 See macro assembler
 5719-CB4
 See COBOL
 5719-F02
 See FORTRAN IV
 5719-LM7
 See macro library
 5719-MS2
 See Multiple Terminal Manager
 5719-PL6
 See PL/I
 5719-SM2
 See Sort/Merge
 5719-SX2
 See \$RJESNA
 5719-UT5
 See utilities
 5719-XS3
 See basic supervisor and emula-
 tor
 5719-XX4
 See program preparation
 5740-LM4
 See macro library/host
 5799-TDE
 See Data Collection Interactive
 (DCI)

7

7850 teletypewriter adapter
 attachment with 3101 Display
 Terminal CT-470
 considerations when
 using CT-80, SG-113
 overview CT-10

READER'S COMMENT FORM

SC34-0404-0

IBM Series/1 Event Driven Executive Indexed Access Method Version 2 Guide

Your comments assist us in improving the usefulness of our publications; they are an important part of the input used in preparing updates to the publications. IBM may use and distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

Please do not use this form for technical questions about the system or for requests for additional publications; this only delays the response. Instead, direct your inquiries or requests to your IBM representative or the IBM branch office serving your locality.

Corrections or clarifications needed:

Page	Comment
------	---------

Cut or Fold Along Line

Please indicate your name and address in the space below if you wish a reply.

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.)

Reader's Comment Form

Cut Along Line

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, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE



IBM Corporation
Information Development, Dept 27T
P.O. Box 1328
Boca Raton, Florida 33432

Fold and tape

Please Do Not Staple

Fold and tape



