

digital

VAX-11
Record Management Services
Reference Manual

Order No. AA-D031C-TE

VAX11

March 1980

This document describes the VAX-11 Record Management Services (RMS). It provides detailed information on the use of VAX-11 RMS facilities with the VAX/VMS operating system.

VAX-11

Record Management Services Reference Manual

Order No. AA-D031C-TE

SUPERSESSION/UPDATE INFORMATION: This document supersedes the VAX-11 Record Management Services Reference Manual (Order No. AA-D031B-TE).

OPERATING SYSTEM AND VERSION: VAX/VMS V02

SOFTWARE VERSION: VAX/VMS V02

To order additional copies of this document, contact the Software Distribution Center, Digital Equipment Corporation, Maynard, Massachusetts 01754

First Printing, August 1978
Revised, February 1979
Revised, March 1980

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may only be used or copied in accordance with the terms of such license.

No responsibility is assumed for the use or reliability of software on equipment that is not supplied by DIGITAL or its affiliated companies.

Copyright © 1978, 1979, 1980 by Digital Equipment Corporation

The postage prepaid READER'S COMMENTS form on the last page of this document requests the user's critical evaluation to assist us in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

DIGITAL	DECsystem-10	MASSBUS
DEC	DECTape	OMNIBUS
PDP	DIBOL	OS/8
DECUS	EDUSYSTEM	PHA
UNIBUS	FLIP CHIP	RSTS
COMPUTER LABS	FOCAL	RSX
COMTEX	INDAC	TYPESET-8
DDT	LAB-8	TYPESET-11
DECCOMM	DECSYSTEM-20	TMS-11
ASSIST-11	RTS-8	ITPS-10
VAX	VMS	SBI
DECnet	IAS	PDT
DATATRIEVE	TRAX	

CONTENTS

			Page
PREFACE			
CHAPTER	1	WHAT IS VAX-11 RMS?	1-1
	1.1	VAX-11 RMS FUNCTIONS	1-1
	1.1.1	Allocating and Initializing Control Blocks	1-1
	1.1.2	Accessing Fields in Control Blocks	1-2
	1.1.3	Requesting File and Record Operations	1-2
	1.2	WHO USES VAX-11 RMS	1-2
	1.3	DEFINITION OF TERMS	1-2
CHAPTER	2	STATEMENT CONVENTIONS	2-1
CHAPTER	3	THE PROGRAM INTERFACE WITH VAX-11 RMS	3-1
	3.1	USER CONTROL BLOCKS	3-1
	3.2	VAX-11 RMS RUN-TIME OPERATIONS	3-2
CHAPTER	4	THE FILE ACCESS BLOCK	4-1
	4.1	THE PURPOSE OF THE FILE ACCESS BLOCK	4-1
	4.2	FAB ALLOCATION	4-3
	4.2.1	Label	4-4
	4.2.2	Allocation Quantity	4-4
	4.2.3	Bucket Size	4-5
	4.2.4	Block Size	4-8
	4.2.5	User Context	4-8
	4.2.6	Default File Extension Quantity	4-9
	4.2.7	Default File Specification String Address	4-10
	4.2.8	Default File Specification String Size	4-10
	4.2.9	Default File Specification	4-11
	4.2.10	File Access	4-11
	4.2.11	File Specification String Address	4-13
	4.2.12	File Specification String Size	4-14
	4.2.13	File Specification	4-14
	4.2.14	File Process Options	4-14
	4.2.15	Fixed Control Area Size	4-18
	4.2.16	Maximum Record Number	4-19
	4.2.17	Maximum Record Size	4-19
	4.2.18	Name Block Address	4-20
	4.2.19	File Organization	4-21
	4.2.20	Record Attributes	4-21
	4.2.21	Record Format	4-23
	4.2.22	Retrieval Window Size	4-24
	4.2.23	File Sharing	4-24
	4.2.24	Extended Attribute Block Pointer	4-26
	4.3	NONINITIALIZABLE FAB FIELDS	4-27

CONTENTS

	Page
CHAPTER 5 THE RECORD ACCESS BLOCK	5-1
5.1 THE PURPOSE OF THE RECORD ACCESS BLOCK	5-1
5.2 RAB ALLOCATION	5-3
5.2.1 Label	5-4
5.2.2 Bucket Code	5-4
5.2.3 Context	5-5
5.2.4 File Access Block Address	5-5
5.2.5 Key Buffer Address	5-5
5.2.6 Key of Reference	5-6
5.2.7 Key Size	5-7
5.2.7.1 Relative Files	5-7
5.2.7.2 Indexed Files	5-7
5.2.8 Multiblock Count	5-8
5.2.9 Multibuffer Count	5-9
5.2.10 Prompt Buffer Address	5-10
5.2.11 Prompt Buffer Size	5-11
5.2.12 Record Access Mode	5-11
5.2.13 Record Address	5-12
5.2.14 Record Header Buffer	5-13
5.2.15 Record-Processing Options	5-13
5.2.16 Record Size	5-18
5.2.17 Time-Out Period	5-19
5.2.18 User Record Area Address	5-20
5.2.19 User Record Area Size	5-20
5.3 NONINITIALIZABLE RAB FIELDS	5-21
5.3.1 The Record's File Address	5-21
CHAPTER 6 THE EXTENDED ATTRIBUTE BLOCKS	6-1
6.1 THE PURPOSE OF EXTENDED ATTRIBUTE BLOCKS	6-1
6.2 CHAINING EXTENDED ATTRIBUTE BLOCKS	6-3
6.3 DATE AND TIME XAB	6-4
6.3.1 Expiration Date and Time	6-5
6.3.2 Creation/Revision Date and Time, and Revision Number	6-6
6.4 FILE PROTECTION XAB	6-7
6.4.1 File Protection	6-8
6.4.2 Group and Member Number	6-10
6.5 ALLOCATION CONTROL XAB	6-11
6.5.1 Area Identification Number	6-12
6.5.2 Alignment Boundary Type	6-13
6.5.3 Allocation Quantity	6-14
6.5.4 Allocation Option	6-14
6.5.5 Bucket Size	6-15
6.5.6 Default Extension Quantity	6-16
6.5.7 Location	6-16
6.5.8 Relative File Identifier	6-17
6.5.9 Relative Volume Number	6-18
6.6 KEY DEFINITION XAB	6-18
6.6.1 Data Bucket Area Number	6-20
6.6.2 Data Bucket Fill Size	6-21
6.6.3 Key Data Type	6-22
6.6.4 Key Options Flag	6-24
6.6.5 Index Bucket Area Number	6-26
6.6.6 Index Bucket Fill Size	6-27
6.6.7 Key Name Address	6-28

CONTENTS

		Page
	6.6.8 Lowest Level of Index Area Number	6-28
	6.6.9 Null Key Value	6-29
	6.6.10 Key Position	6-30
	6.6.11 Key of Reference	6-31
	6.6.12 Key Size	6-32
	6.6.13 Noninitializable Key Fields	6-33
	6.7 SUMMARY XAB	6-34
	6.8 FILE HEADER CHARACTERISTICS XAB	6-35
	6.9 REVISION DATE AND TIME XAB	6-37
	6.9.1 Revision Date and Time	6-38
CHAPTER	7 THE NAME BLOCK	7-1
	7.1 THE PURPOSE OF THE NAME BLOCK	7-1
	7.2 NAM BLOCK ALLOCATION	7-2
	7.2.1 Label	7-3
	7.2.2 Expanded String Area Address	7-3
	7.2.3 Expanded String Area Size	7-4
	7.2.4 Related File Nam Block Address	7-4
	7.2.5 Resultant String Area Address	7-5
	7.2.6 Resultant String Area Size	7-5
	7.3 NONINITIALIZABLE NAM BLOCK FIELDS	7-6
CHAPTER	8 RUN-TIME PROCESSING INTERFACE	8-1
	8.1 THE VAX-11 RMS CALLING SEQUENCE	8-1
	8.2 THE PATH TO A FILE	8-3
	8.2.1 Interpretation of the File Specification	8-4
	8.2.2 Wild Card Characters in File Specifications	8-5
	8.2.3 File Specification Default Application	8-6
	8.2.4 Opening and Creating a File by Name Block	8-7
	8.3 CONTROL BLOCK USAGE	8-7
	8.4 COMPLETION STATUS CODES	8-8
	8.5 PROCESS PERMANENT FILES	8-9
CHAPTER	9 FILE-PROCESSING MACRO INSTRUCTIONS	9-1
	9.1 TERMINATING FILE PROCESSING	9-1
	9.2 CREATING A FILE	9-4
	9.3 OBTAINING ATTRIBUTES OF A FILE	9-8
	9.4 DELETING A FILE	9-10
	9.5 EXTENDING A FILE'S ALLOCATED SPACE	9-13
	9.6 OPENING AN EXISTING FILE	9-14
CHAPTER	10 RECORD OPERATION PERFORMANCE	10-1
	10.1 RECORD ACCESS	10-1
	10.1.1 Specifying the Record Access Mode	10-1
	10.1.2 Specifying the Record Transfer Mode	10-2
	10.2 CURRENT RECORD CONTEXT	10-3
	10.2.1 Current Record	10-3
	10.2.2 Next Record	10-4
	10.3 RECORD STREAMS	10-5
	10.4 SYNCHRONOUS AND ASYNCHRONOUS OPERATIONS	10-7
	10.4.1 Synchronous Operations	10-7
	10.4.2 Asynchronous Operations	10-8

CONTENTS

		Page
	10.5 FILE SHARING	10-8
	10.6 RECORD LOCKING	10-9
	10.6.1 Automatic Record Locking	10-10
	10.6.2 Manual Record Locking	10-11
	10.6.3 Controlling Record Locking	10-12
CHAPTER	11 RECORD-PROCESSING MACRO INSTRUCTIONS	11-1
	11.1 ESTABLISHING A RECORD STREAM	11-2
	11.2 DELETING A RECORD	11-3
	11.3 TERMINATING A RECORD STREAM	11-5
	11.4 LOCATING A RECORD	11-7
	11.5 WRITING OUT MODIFIED I/O BUFFERS	11-9
	11.6 UNLOCKING ALL RECORDS	11-11
	11.7 RETRIEVING A RECORD	11-12
	11.8 CONTINUE PROCESSING ON NEXT VOLUME	11-17
	11.9 WRITING A RECORD TO A FILE	11-19
	11.10 UNLOCKING A RECORD	11-22
	11.11 POSITIONING TO THE FIRST RECORD	11-24
	11.12 TRUNCATING A SEQUENTIAL FILE	11-26
	11.13 UPDATING AN EXISTING RECORD	11-27
	11.14 STALL FOR I/O COMPLETION	11-30
CHAPTER	12 PERFORMING BLOCK I/O	12-1
	12.1 TRANSFER TO MEMORY	12-3
	12.2 POSITIONING TO A BLOCK	12-5
	12.3 WRITE TO A FILE	12-7
	12.4 NON-FILE-STRUCTURED OPERATIONS	12-9
CHAPTER	13 FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS	13-1
	13.1 ENTER A FILE NAME	13-1
	13.2 PARSE A FILE SPECIFICATION STRING	13-4
	13.3 REMOVE A FILE NAME	13-6
	13.4 RENAME A FILE	13-8
	13.5 SEARCH FOR FILE NAME	13-11
CHAPTER	14 RUN-TIME CONTROL BLOCK INITIALIZATION	14-1
	14.1 THE STORE MACRO INSTRUCTIONS	14-1
CHAPTER	15 CONTROL ROUTINES	15-1
	15.1 HALT I/O AND CLOSE FILES	15-1
	15.2 SET DEFAULT DIRECTORY	15-2
	15.3 SET DEFAULT FILE PROTECTION	15-2
APPENDIX A	COMPLETION STATUS CODES	A-1
APPENDIX B	FILE/RECORD CONCEPTS AND FORMATS	B-1
	B.1 FILE ORGANIZATIONS	B-1
	B.2 RECORD ACCESS MODES	B-2
	B.3 RECORD FORMATS	B-4
	B.4 FILES-11 DISK STRUCTURE	B-4
	B.4.1 Files-11 Directories	B-9

CONTENTS

		Page
B.5	MAGNETIC TAPE HANDLING	B-10
B.5.1	Volume Label	B-11
B.5.2	File Header Label	B-14
B.5.3	End-of-file And End-of-volume Labels	B-19
B.5.4	Arrangement of Labels And Data	B-20
APPENDIX C	FILE SPECIFICATION PARSING	C-1
C.1	LOGICAL-NAME-OR-FILE-NAME SYNTAX	C-2
C.2	QUOTED-STRING-SPECIFICATION SYNTAX	C-3
C.3	FULL-FILE-SPECIFICATION SYNTAX	C-4
APPENDIX D	DIGITAL-ONLY COMPONENT OPTIONS	D-1
INDEX		Index-1

FIGURES

FIGURE	6-1	File Protection Field	6-8
	8-1	Argument List Format	8-2
	B-1	Logical and Virtual Block Numbers	B-6
	B-2	Volume Label Format	B-11
	B-3	HDR1 Label Format	B-14
	B-4	HDR2 Label Format	B-17
	B-5	Single File, Single Volume	B-20
	B-6	Single File, Multivolume	B-20
	B-7	Multifile, Single Volume	B-21
	B-8	Multifile, Multivolume	B-21

TABLES

TABLE	3-1	User Control Blocks	3-2
	3-2	Run-Time Processing Macro Instructions	3-3
	4-1	File Access Block Fields	4-2
	4-2	Device Characteristics	4-29
	5-1	Record Access Block Fields	5-2
	6-1	XAB Types Processed by Service	6-3
	6-2	Date and Time Extended Attribute Block Fields	6-5
	6-3	File Protection Extended Attribute Block Fields	6-7
	6-4	Allocation Control Extended Attribute Block Fields	6-11
	6-5	Key Definition Extended Attribute Block Fields	6-19
	6-6	Key Field Data Types, Data Type Codes and Global Symbols	6-22
	6-7	Packed Decimal Digits and Signs Representation	6-23
	6-8	Key Options Flag Combinations	6-25
	6-9	Summary Extended Attribute Block Fields	6-34
	6-10	File Header Characteristics Extended Attribute Block Fields	6-36

CONTENTS

TABLES (Cont.)

			Page
TABLE	6-11	Revision Date and Time Extended Attribute	
		Block Fields	6-37
	7-1	Name Block Fields	7-2
	7-2	File Name Status Bits	7-8
	9-1	Close FAB Fields	9-3
	9-2	Create FAB Fields	9-6
	9-3	Create NAM Block Fields	9-7
	9-4	Display FAB Fields	9-9
	9-5	Erase FAB Fields	9-11
	9-6	Erase NAM Block Fields	9-12
	9-7	Extend FAB Fields	9-14
	9-8	Open FAB Fields	9-16
	9-9	Open NAM Block Fields	9-17
	10-1	Record Access Stream Context	10-6
	11-1	Connect RAB Fields	11-3
	11-2	Delete RAB Fields	11-5
	11-3	Disconnect RAB Fields	11-6
	11-4	Find RAB Fields	11-8
	11-5	Flush RAB Fields	11-10
	11-6	Free RAB Fields	11-12
	11-7	Get RAB Fields	11-15
	11-8	Next Volume RAB Fields	11-18
	11-9	Put RAB Fields	11-21
	11-10	Release RAB Fields	11-23
	11-11	Rewind RAB Fields	11-25
	11-12	Truncate RAB Fields	11-27
	11-13	Update RAB Fields	11-29
	11-14	Wait RAB Fields	11-31
	12-1	Read RAB Fields	12-4
	12-2	Space RAB Fields	12-6
	12-3	Write RAB Fields	12-8
	13-1	Enter Fields	13-3
	13-2	Parse Fields	13-5
	13-3	Remove Fields	13-7
	13-4	Rename Fields	13-10
	13-5	Search Fields	13-13
	B-1	File Organization Relationships with Record Access Modes and Record Formats	B-5
	B-2	Search Delta Geometry	B-7
	B-3	Volume Label Contents	B-12
	B-4	HDR1 Label Contents	B-15
	B-5	HDR2 Label Contents	B-18
	B-6	HDR3 Label Contents	B-19

PREFACE

MANUAL OBJECTIVES

The intent of this manual is to enable VAX-11 MACRO programmers to use the VAX-11 Record Management Services (RMS) facilities provided by the VAX/VMS operating system.

Many data operations can be performed by using VAX-11 RMS and associated control routines. You can perform these operations by simply calling a VAX-11 RMS routine with the appropriate parameters, rather than writing your own routines.

INTENDED AUDIENCE

VAX/VMS provides record management services for all the supported languages. Except for VAX-11 MACRO, each particular language manual provides the necessary information about performing record management. However, for the VAX-11 MACRO programmers, and for those high-level language programmers who wish to call VAX-11 RMS directly, this manual contains a description of the user interface to record management.

STRUCTURE OF THIS DOCUMENT

This manual consists of three parts, as follows:

Part I: Introduction to VAX-11 RMS

Part I, consisting of Chapters 1 and 2, discusses VAX-11 RMS in terms of who uses it and why.

Part II: VAX-11 RMS Program Interface

Part II can be subdivided in the following way. Chapters 3 through 7 describe the fields for VAX-11 RMS structures, such as file declaration and the macro instructions used to initialize these fields. Chapters 8 through 15 describe the interfaces to VAX-11 RMS file and record operations and control routines.

Appendixes

The appendixes summarize the concepts of files and records, provide formulas for determining file and record size, and list completion status codes.

ASSOCIATED DOCUMENTS

The following manuals are related to this document:

- Introduction to VAX-11 Record Management Services
- VAX-11 Record Management Services User's Guide
- RMS-11 User's Guide
- VAX-11 MACRO User's Guide
- VAX-11 MACRO Language Reference Manual
- VAX/VMS System Services Reference Manual

The Introduction to VAX-11 Record Management Services manual contains introductory information about file services and structures in general, and about VAX-11 RMS in particular. The VAX-11 RMS User's Guide contains detailed information on using the capabilities of VAX-11 RMS efficiently. Much of this information is illustrated in programming examples. The RMS-11 User's Guide also contains useful information concerning file processing, much of which is compatible with VAX-11 RMS.

For a complete list of all VAX-11 documents, including brief descriptions of each, see the VAX-11 Information Directory and Index.

SUMMARY OF TECHNICAL CHANGES

This manual has been revised to reflect VAX-11 RMS support for wild card characters and for file sharing for sequential files with 512-byte fixed-length records.

CHAPTER 1

WHAT IS VAX-11 RMS?

The VAX-11 Record Management Services (VAX-11 RMS) are generalized routines that assist user programs in processing and managing files and their contents. VAX-11 RMS also includes a set of macro instructions that you can use to initialize control blocks and call VAX-11 RMS service routines.

1.1 VAX-11 RMS FUNCTIONS

VAX-11 RMS provides a variety of file organizations and record access modes that let you choose the processing techniques best suited to your application. VAX-11 RMS organizes files sequentially, relatively, or in indexed form. You can access records in these files in a number of ways:

- Sequentially
- Randomly by key
- Randomly by the record's file address (RFA)
- Dynamically, which is a mixture of sequential and random access

You transmit file and record operation requests to VAX-11 RMS through control blocks. Through these same control blocks, such as the File Access Block or Record Access Block, VAX-11 RMS returns to you the data contents of files, attribute information about the files, and status codes.

To use VAX-11 RMS, you must:

- Allocate and initialize control blocks
- Access fields in these control blocks at run time
- Request a particular file or record operation through the use of macro instructions

1.1.1 Allocating and Initializing Control Blocks

You communicate with VAX-11 RMS through control blocks. You must allocate space in your program for the control blocks; usually, this is done at assembly time. In addition, you can establish initial values for the fields in these blocks through assembly-time initialization macros.

WHAT IS VAX-11 RMS?

1.1.2 Accessing Fields in Control Blocks

At run time, you can store values in the control block data fields through the use of macro instructions, or you can access data in the control block fields directly by using the defined offsets for the fields.

1.1.3 Requesting File and Record Operations

Control blocks combined with a set of VAX-11 RMS file and record operation macro instructions form the complete run-time program interface with VAX-11 RMS. Each macro instruction represents a request for a particular VAX-11 RMS file or record service. The fields of the control blocks further describe the request. Using VAX-11 RMS macro instructions, you can:

- Create new files
- Process existing files
- Extend and delete files
- Read, write, update, and delete records within files

1.2 WHO USES VAX-11 RMS

VAX-11 MACRO programmers make direct use of the VAX-11 RMS routines. Programmers writing in a high-level language, such as VAX-11 FORTRAN, can write their programs to interface with VAX-11 RMS facilities either 1) directly through the use of a call facility in the language, or 2) indirectly through the input/output (I/O) instructions of the language. The latter interface is much more commonly used. Programs that interface directly with VAX-11 RMS can use all its capabilities, whereas programs that use an I/O statement of a high-level language are generally restricted to the subset of VAX-11 RMS capabilities used by that language. This manual, describing the full VAX-11 RMS interface, is therefore directed primarily to the VAX-11 MACRO user. High-level language users should see the VAX-11 manuals specific to their language.

1.3 DEFINITION OF TERMS

The following glossary defines terms that appear throughout this manual.

alternate key

An optional key within the data records in an indexed file; used by VAX-11 RMS to build an alternate index. See key (indexed files) and primary key.

area

VAX-11 RMS-maintained region of an indexed file which are used for allocating buckets. An area consists of any number of buckets, and there may be from 1 through 255 areas in a file.

WHAT IS VAX-11 RMS?

block

A unit of I/O transfer. A block on a Files-11 disk structure is fixed at 512 bytes and contains one or more complete or partial records. A block on tape contains one or more complete records; its size is user determined.

block I/O

An I/O technique using a set of VAX-11 RMS procedures that allow direct access to the blocks in a file, regardless of the file organization or record format.

bucket

A structure used to store and transfer blocks of data for a relative or indexed file. A bucket consists of from 1 through 32 blocks.

buffer

An area in memory used to store data temporarily during input or output operations.

cluster

The basic unit of space allocation on a Files-11 disk. A cluster consists of one or more blocks, as defined by the initializer of the disk.

directory name

The field in a file specification that identifies the directory in which the file is listed. It begins with a left bracket ([or <) and ends with a right bracket (] or >). The brackets enclose either a group number and a user number separated by a comma, or an alphanumeric directory list.

dynamic access

The process of switching from one record access mode to another while processing a file.

extent

One or more adjacent clusters allocated to a file or a portion of a file.

file

A collection of data; generally used to refer to data stored on a magnetic medium, such as a disk.

file header

A block in the index file that describes a file on a Files-11 disk. Every file residing on the disk has at least one file header, which provides the location of the file's extents.

file organization

The physical arrangement of data in a file. VAX-11 RMS supports three file organizations -- sequential, relative, and indexed.

file specification

The alphanumeric character string that specifies a file within the system.

Files-11

The standard VAX-11 RMS physical disk structure.

WHAT IS VAX-11 RMS?

fixed control area

An area, prefixed to a variable-length record, containing additional information about the record that may have no bearing on the other contents of the record. For example, the fixed control area may contain line numbering or carriage control information.

fixed-length record format

The property of a file specifying that all records must be the same length. This format allows for simplicity in determining the exact location of a record in the file and eliminates the need to prefix a record size field to each record.

home block

A block in the volume's index file that contains information pertaining to the volume as a whole, such as volume label and protection.

index

The structure which allows retrieval by key value of records in an indexed file. See key (indexed files).

index file

The file on a Files-11 volume that provides the means for identification and initial access to the volume. The index file contains the access data for all files on the volume (including itself).

indexed file organization

A file organization which allows random retrieval of records by key value and sequential retrieval of records within the key of reference. See key (indexed files).

key

indexed files: A character string, a packed decimal number, a 2- or 4-byte unsigned binary number, or a 2- or 4-byte signed integer within each data record in an indexed file; it is user defined as to length and location within the records; VAX-11 RMS uses the key to build an index. See primary key, alternate key, and random access by key (indexed files only).

key

relative files: The relative record number of each data record in a data file; VAX-11 RMS uses the relative record numbers to identify and access data records in a relative file in random access mode. See relative record number.

locate mode

Record transfer technique in which records stay in place while operations are performed. The records are not copied from the I/O buffer to a user buffer; the address of the record in the I/O buffer is returned to the user.

logical block number

The number assigned to a block on a disk volume, sequentially beginning with 0 through the number of blocks that will fit on the volume. See also virtual block number.

move mode

Record transfer technique in which a record is copied between an I/O buffer and a user buffer.

WHAT IS VAX-11 RMS?

primary key

The mandatory key within the data records of an indexed file; used by VAX-11 RMS to build a primary index; see key (indexed files) and alternate key.

process permanent file

A file opened or created through VAX-11 RMS in supervisor or executive mode. The internal data structures of a process permanent file are allocated such that the file may be open across image activations; a restricted subset of allowable operations is available to "indirect" accessors.

random access by key

for indexed files: Retrieval of a data record in an indexed file by the primary (or optionally, alternate) key within the data record. See key (indexed files).

for relative files or sequential files with 512-byte fixed-length records: Retrieval of a data record in a relative file by the relative record number of the record. See key (relative files).

random access by record's file address

The retrieval of a record by the record's unique address that VAX-11 RMS returns to the user. This record access mode is the only means of randomly accessing a sequential file containing variable-length records.

random access by relative record number

The retrieval of a record by specifying the record's number relative to the beginning of the file. For relative files, random access by relative record number is synonymous with random access by key. See random access by key (relative files only).

record

A collection of related data within a file treated as a unit of information.

record access mode

The manner in which VAX-11 RMS selects the next record to be accessed, that is, sequentially or randomly.

record cell

A fixed-length area in a relative file that is capable of containing a record. The concept of a fixed-length record cell lets VAX-11 RMS make a direct calculation of the record's actual position in the file.

record's file address

The unique address of a record in a file. This address allows records to be accessed randomly regardless of file organization. It is valid only for a particular instance of a file.

record format

The way a record physically appears on the recording surface of the storage medium. The record format defines the method for determining record length.

record locking

A facility that prevents concurrent access to a record by more than one record stream or process until the initiating record stream or process releases the record.

record length

The size of a record, expressed as a number of bytes.

WHAT IS VAX-11 RMS?

relative file organization

The arrangement of records in a file where each record occupies a cell of equal length within a bucket. Each cell is assigned a successive number, which represents its position relative to the beginning of the file.

relative record number

An identification number that specifies the position of a record cell relative to the beginning of the file; used as the key during random access by key mode to relative files.

RFA

See Record's File Address

sequential file organization

The arrangement of records in a file in a sequential fashion. Records appear in the order in which they were written.

sequential record access mode

The retrieval or storage of records starting at a designated point in the file and continuing to access additional records in the order in which they logically appear.

spooling

The technique of using a high-speed mass storage device (such as a disk) to buffer data passing between high-speed main memory and low-speed I/O devices (such as line printers). The high-speed mass storage device (the intermediate device) temporarily stores the data passing to and from the low-speed device (the spooled device). The data is queued on the intermediate device to await transmission to the printer for printing (output spooling) or to the processor for processing (input spooling).

storage allocation

The assignment of space to a file on the recording medium.

user identification code

The number assigned to a user identifying the user and, consequently, determining the files to which the user has access. It consists of a group number and a user number, separated by a comma, and enclosed in brackets, i.e., [100,5].

variable-length record format

The property of a file specifying that records need not be the same length.

variable with fixed-length control record format

The property of a file specifying that records of variable-length contain an additional fixed control area capable of storing data that may have no bearing on the other contents of the record. Variable with fixed-length control record format is not applicable to indexed files.

VAX-11 Record Management Services (VAX-11 RMS)

The file and record access system for the VAX/VMS operating system. VAX-11 RMS allows programs to issue requests at the record and block level.

virtual block number

The number assigned to a block of a file. This number refers to the position of the block relative to other blocks in the same file, instead of to its position relative to other blocks on the volume. Virtual block numbers are assigned to the blocks of a file beginning with 1. The file header provides relocation information for mapping the file's virtual block numbers to the volume's logical block numbers. See also logical block number.

CHAPTER 2

STATEMENT CONVENTIONS

Throughout this manual, certain conventions apply to the syntax of the VAX-11 RMS macro instructions and control routines.

In examples, parameters other than the parameter under discussion are shown. The purpose of showing these additional parameters is to illustrate and reconfirm throughout the manual some of the conventions that apply in coding macro instructions, such as statement continuation and parameter separation. The parameter under discussion will be shown in red print.

For example:

```
$FAB FNA=FLNAM ALQ=132 BKS=4
```

In coding VAX-11 RMS macro instructions, you follow the same coding rules used by the VAX-11 MACRO assembler. These rules are repeated below for ease of reference.

- Comments must be separated from the rest of the code line by a semicolon (;). For example:

```
$FAB BKS=4 ;bucket size
```

- All the parameters necessary for a macro instruction must be coded on a single macro instruction. If the parameters needed do not all fit on one line (or if you do not want them on one line), you can type the continuation character -- hyphen (-) -- as the last character on the line, and then continue typing parameters on the next line. Comments can follow the hyphen, separated by the comment-delimiting semicolon -- they are not interpreted as code. For example:

```
$FAB FNA=FLNAM - ; filename address
      ALQ=132 - ; allocation quantity
      BKS=4 ; bucket size
```

- Parameters and subparameters can be separated from each other by:
 - A single comma, with or without spaces or tabs; the preferred usage is the comma without a space or tab. That is how coding examples appear in this manual.

```
FNA=FLNAM,ALQ=132
```

STATEMENT CONVENTIONS

- A blank space

FNA=FLNAM ALQ=132

- Multiple blank spaces or tabs

FNA=FLNAM ALQ=132

- Lowercase letters and words represent information that you must supply. Such lowercase information may contain hyphens for readability. The accompanying text defines the information to be supplied. For example:

window-size
address

- Uppercase letters and words, equal signs (=), angle brackets (<>), and dollar signs (\$), must be coded as shown. For example:

RAT=<BLK,CR>
\$OPEN

- Information enclosed within braces indicates that you may choose any one of the enclosed values. For example:

FIX
VAR
VFC
UDF

- Each option has its own symbolic bit offset and mask value. The bit offset is formed by prefixing the control block name and \$V_ to the option value. For example:

FAB\$V_PUT
RAB\$V_ASY

The mask value is formed by prefixing the control block name and \$M_ to the option value. For example:

FAB\$M_PUT
RAB\$M_ASY

CHAPTER 3

THE PROGRAM INTERFACE WITH VAX-11 RMS

You gain access to the VAX-11 RMS facilities at run time by calling record management services. Your program and VAX-11 RMS exchange information by means of user control blocks defined within your program. This chapter provides an introduction to these services and user control blocks, and the macro instructions that facilitate their use.

With each request for a VAX-11 RMS service, you must place the information detailing this request in a user control block. For example, a request to open a file must be accompanied by the name of the file, information on sharing the file, and details on accessing the file. Or, as another example, a program request to read a record from a file must specify a record access mode, or perhaps a buffer size.

Once a request for a service is satisfied, VAX-11 RMS uses the same user control block to return information to your program. For example, when the file is successfully opened, VAX-11 RMS returns attribute information, such as file organization and record format. Or, when a record is retrieved from a file, VAX-11 RMS provides your program with the record's length and location in memory.

The amount of information exchanged between VAX-11 RMS and your program varies with the nature of the request and the file attributes.

The following sections provide a broad overview of the interface that a program uses when requesting VAX-11 RMS services. The remaining chapters of Part II present detailed information on using the VAX-11 RMS declarative and imperative macro instructions. The declarative macro instructions allocate and initialize file access blocks (FABs), record access blocks (RABs), name blocks (NAMs), and extended attribute blocks (XABs). The imperative macro instructions invoke VAX-11 RMS operations to manipulate files and records.

3.1 USER CONTROL BLOCKS

You must allocate user control blocks as formatted areas in your program. Your program and VAX-11 RMS use the data fields in these blocks to exchange information.

Usually, you allocate space for user control blocks at assembly time. Optionally, you can also set values for the fields in these blocks either initially or at run time. The VAX-11 RMS declarative macro instructions perform the functions that support assembly-time allocation and initialization. For efficiency, align the control blocks on a longword boundary; if you do not, you will receive a warning message from the assembler. Since VAX-11 RMS returns

THE PROGRAM INTERFACE WITH VAX-11 RMS

information in the fields of these user control blocks, you cannot allocate user control blocks in read-only storage.

Table 3-1 lists the user control blocks that are part of your program interface with VAX-11 RMS. The Macro Name column shows the VAX-11 RMS macro instruction you use to allocate space for the control block. Chapters 4 through 7 describe these macro instructions.

Table 3-1
User Control Blocks

Block Name	Function	Macro Name
File Access Block FAB	Describes a file and contains file-related information	SFAB
Record Access Block RAB	Describes a record and contains record-related information	SRAB
Extended Attribute Blocks XAB	Contains file attribute information beyond that in the File Access Block	\$XABxxx ¹
Name Block NAM	Contains file specification information beyond that in the File Access Block	SNAM

¹ xxx is a 3-character XAB type specification.

3.2 VAX-11 RMS RUN-TIME OPERATIONS

To create and process VAX-11 RMS files, your program must contain calls to appropriate VAX-11 RMS routines. Generally, you make these calls by using the VAX-11 RMS imperative macro instructions for run-time processing. The expanded code of these macro instructions, when encountered at run time, causes calls to be made to the corresponding VAX-11 RMS routine. Each macro instruction, and the resultant call, represents a program request for either a file or record related service, or block I/O transfer operation.

Table 3-2 summarizes the run-time processing macro instructions. Chapters 8 through 15 describe these macro instructions.

THE PROGRAM INTERFACE WITH VAX-11 RMS

**Table 3-2
Run-Time Processing Macro Instructions**

Category	Macro Name	Service
File Processing	\$CREATE	Creates and opens a new file of any organization
	\$OPEN	Opens an existing file and initiates file processing
	\$DISPLAY	Returns the attributes of a file to user program
	\$EXTEND	Extends the allocated space of a file
	\$CLOSE	Terminates file processing and closes the file
	\$ERASE	Deletes a file and removes its directory entry
Record Processing	\$GET	Retrieves a record from a file
	\$PUT	Writes a new record to a file
	\$UPDATE	Rewrites an existing record in a file
	\$DELETE	Deletes a record from a relative or indexed file
	\$FIND	Locates and positions to a record and returns its RFA
	\$CONNECT	Associates and connects a RAB to a file
	\$DISCONNECT	Disconnects a RAB from a file
	\$RELEASE	Unlocks a record pointed to by the contents of the RFA field of the RAB
	\$FREE	Unlocks all previously locked records
	\$WAIT	Determines the completion of an asynchronous record operation
	\$REWIND	Positions to the first record of a file
	\$TRUNCATE	Truncates a sequential file
	\$FLUSH	Write modified I/O buffers and file attributes
	\$NXTVOL	Causes processing of a magnetic tape file to continue to the next volume of a volume set
Block I/O	\$READ	Retrieves a specified number of bytes from a file
	\$WRITE	Writes a specified number of bytes to a file
	\$SPACE	Spaces forward or backward in a file
File Naming	\$ENTER	Enters a file name into a directory
	\$PARSE	Parses a file specification
	\$REMOVE	Removes a file name from a directory
	\$RENAME	Assigns a new name to a file
	\$SEARCH	Searches a directory for a file name

CHAPTER 4

THE FILE ACCESS BLOCK

This chapter describes the File Access Block (FAB), the fields in the FAB, and the parameters of the \$FAB macro instruction. The FAB is used by the file processing services (Chapter 9) and the file specification processing services (Chapter 13).

4.1 THE PURPOSE OF THE FILE ACCESS BLOCK

The FAB is a user control block that describes a particular file. The fields of the FAB contain file-related information, such as:

- The name of the file
- The file organization
- The record format
- Disk storage space allocation information

You allocate a FAB with a \$FAB macro instruction, and initialize the fields of the FAB either at assembly time (through keyword parameters) or by direct manipulation at run time. You initialize the FAB at run time through either keyword parameters with the \$FAB STORE macro instruction (see Chapter 14) or the defined symbolic offsets. You need one FAB for each open file in your program.

Each field in the FAB has a 3-character mnemonic name. All access to these fields is through this name (by keyword or offset). However, some of the fields are static or output-only; therefore, you need not initialize them. Table 4-1 summarizes the fields of the FAB, including the static and output-only fields.

THE FILE ACCESS BLOCK

Table 4-1
File Access Block Fields

Field & Keyword Name	Field Size (units of 1)	Description	Offset
ALQ	longword	Allocation quantity	FAB\$ <u>L</u> _ALQ
BID ¹	byte	Block identifier	FAB\$ <u>B</u> _BID
BKS	byte	Bucket size	FAB\$ <u>B</u> _BKS
BLN ¹	byte	Block length	FAB\$ <u>B</u> _BLN
BLS	word	Block size	FAB\$ <u>W</u> _BLS
CTX	longword	Context	FAB\$ <u>L</u> _CTX
DEQ	word	Default file extension quantity	FAB\$ <u>W</u> _DEQ
DEV ²	longword	Device characteristics	FAB\$ <u>L</u> _DEV
DNA	longword	Default file specification string address	FAB\$ <u>L</u> _DNA
DNS	byte	Default file specification string size	FAB\$ <u>B</u> _DNS
FAC	byte	File access	FAB\$ <u>B</u> _FAC
FNA	longword	File specification string address	FAB\$ <u>L</u> _FNA
FNS	byte	File specification string size	FAB\$ <u>B</u> _FNS
FOP	longword	File-processing options	FAB\$ <u>L</u> _FOP
FSZ	byte	Fixed control area size	FAB\$ <u>B</u> _FSZ
IFI ²	word	Internal file identifier	FAB\$ <u>W</u> _IFI
MRN	longword	Maximum record number	FAB\$ <u>L</u> _MRN
MRS	word	Maximum record size	FAB\$ <u>W</u> _MRS
NAM	longword	Name block address	FAB\$ <u>L</u> _NAM
ORG	byte	File organization	FAB\$ <u>B</u> _ORG
RAT	byte	Record attributes	FAB\$ <u>B</u> _RAT
RFM	byte	Record format	FAB\$ <u>B</u> _RFM
RTV	byte	Retrieval window size	FAB\$ <u>B</u> _RTV
SDC ²	longword	Spooling device characteristics	FAB\$ <u>L</u> _SDC
SHR	byte	File sharing	FAB\$ <u>B</u> _SHR
STS ²	longword	Completion status code	FAB\$ <u>L</u> _STS
STV ²	longword	Status values	FAB\$ <u>L</u> _STV
XAB	longword	Extended attribute block address	FAB\$ <u>L</u> _XAB

¹ Indicates statically initialized field (by \$FAB macro instruction) to identify this control block as a FAB.

² Indicates nonuser-initialized field.

\$FAB

4.2 FAB ALLOCATION

The format of the \$FAB macro instruction is shown below. Every parameter is optional, depending on the function to be performed with the FAB and the combination of parameters in the macro instruction as a whole.

Format:

OPERATION	PARAMETERS
label: \$FAB	ALQ=allocation-qty BKS=bucket-size BLS=block-size CTX=value DEQ=extension-qty DNA=address DNM=<filespec> DNS=value FAC=<PUT GET DEL UPD TRN BIO BRO> FNA=address FNM=<filespec> FNS=value FOP=<CBT CIF CTG DFW DLT MXV NAM NEF NFS OFP POS RCK RWC RWO SCF SPL SQO SUP TEF TMD TMP UM UFO WCK> FSZ=header-size MRN=max-rec-number MRS=max-rec-size NAM=nam-address ORG= { REL } { SEQ } { IDX } RAT=<BLK { CR } { FTN }> { PRN } RFM= { FIX } { VAR } { VFC } { UDF } RTV=window-size SHR=<PUT GET DEL UPD NIL MSE UPI> XAB=xab-address

THE FILE ACCESS BLOCK

The \$FAB macro instruction allocates and initializes storage for a FAB. You cannot use this macro instruction within a sequence of executable instructions. In some cases, specific default values are assigned automatically, when you omit a parameter. These specific defaults are noted in the text that explains each parameter. If there is no specific default, VAX-11 RMS uses a default value of 0.

label: \$FAB

4.2.1 Label

You can use the label field of the \$FAB macro instruction to name a FAB and thereby to refer to a particular FAB within your program. The label field is optional but when used, must precede the symbol \$FAB and be separated from \$FAB by a colon (:). For example:

```
INFAB: $FAB
```

\$FAB ALQ

4.2.2 Allocation Quantity

You can use the ALQ parameter to initialize the allocation quantity field. With this field you can specify the amount of space, in blocks, to be initially allocated to a disk file when it is created, or to be added to the file when it is explicitly extended (through a \$EXTEND macro instruction).

Format

```
ALQ=allocation-quantity
```

allocation-quantity

A numeric value representing a number of blocks, in the range of 0 through 4,294,967,295. A value of 0 indicates no allocation.

For example, to set an allocation quantity of 132 blocks, the coding is:

```
$FAB ALQ=132
```

User Considerations

1. When you create a new file with a \$CREATE macro instruction, VAX-11 RMS interprets the value in the allocation quantity field as the number of blocks for the initial extent of the file. If the value is 0, the minimum number of blocks for the specific file organization is the allocation quantity used for the initial extent. For example, in indexed files, the number of blocks necessary to contain key and area definitions is used as the initial extent quantity when ALQ=0.
2. When an existing file is opened with a \$OPEN macro instruction, VAX-11 RMS sets the allocation quantity field to indicate the highest virtual block number currently allocated to the file.

THE FILE ACCESS BLOCK

3. Before extending a file with a \$EXTEND macro instruction, you must set the allocation quantity field equal to the number of blocks to be added to the file. You cannot use an extension size of 0.
4. When you use the \$CREATE and \$EXTEND macro instructions, the allocation quantity value is rounded up to the next cluster boundary; the number of blocks actually allocated is returned in the allocation quantity field.

NOTE

The function of the allocation quantity field with the \$CREATE and \$EXTEND macro instructions is different from the preceding description if allocation XABs are present during the operation. Chapter 6 describes allocation XABs and their effect on the allocation quantity field during file creation or extension.

\$FAB BKS

4.2.3 Bucket Size

The BKS parameter initializes the bucket size field. This field is used only for relative or indexed files. When you open an existing relative or indexed file, VAX-11 RMS sets the bucket size field to the defined size of the buckets in the file. However, when you create a new relative or indexed file, you must set the bucket size field before you issue the \$CREATE macro instruction.

NOTE

If allocation control XABs are specified, the value specified in the XAB BKZ field will supersede the value specified in the FAB BKS field. Refer to Section 6.5.6 for a description of the XAB BKZ parameters.

Format

BKS=bucket-size

bucket-size

A numeric value, in the range of 0 through 32, representing the number of blocks in each bucket of the file. If you omit this parameter or use a value of 0, you receive a default size equal to the minimum number of blocks required to contain a single record.

For example, to set the bucket size to 4, the syntax is:

```
$FAB BKS=4,ALQ=132
```

THE FILE ACCESS BLOCK

User Considerations

In specifying a bucket size, you must be aware of the relationship between bucket size and record size. Since VAX-11 RMS does not allow records to cross bucket boundaries, you must ensure that the number of blocks per bucket conforms to one of the following formulas:

- Relative files with fixed-length records:

$$\text{Bsiz} = ((\text{Rlen}+1)*\text{Rnum})/512$$

where

Bsiz is the number of blocks per bucket rounded up to the next higher integer. The result must be in the range from 1 through 32.

Rlen is the fixed record length.

Rnum is the number of records that you want in each bucket.

- Relative files with variable-length records:

$$\text{Bsiz} = ((\text{Rmax}+3)*\text{Rnum})/512$$

where

Bsiz is the same as described above.

Rmax is the maximum size of any record in the file.

Rnum is the number of records that you want in each bucket. Variable-length records in a relative file bucket always occupy Rmax+3 bytes.

- Relative files with variable with fixed-length control records:

$$\text{Bsiz} = ((\text{Rmax}+\text{Fsiz}+3)*\text{Rnum})/512$$

where

Bsiz is the same as described above.

Rmax is the maximum size of the data portion of any record in the file.

Fsiz is the size of the fixed control area portion of the records.

Rnum is the number of records that you want in each bucket. Variable with fixed-length control records in a relative file bucket always occupy Rmax+Fsiz+3 bytes.

THE FILE ACCESS BLOCK

- Indexed files with fixed-length records:

$$\text{Bsiz} = ((\text{Rlent}+7)*\text{Rnum})+15/512$$

where

Bsiz is the number of blocks per bucket rounded up to the next higher integer. The result must be in the range of from 1 through 32.

Rlen is the fixed-record length.

Rnum is the number of records that you want in each bucket. Fixed-length records in an indexed file bucket always occupy Rlen plus seven bytes of record control information. Fifteen bytes are required for bucket control information.

- Indexed files with variable-length records:

$$\text{Bsiz} = ((\text{Rmax}+9)*\text{Rnum})+15/512$$

where

Bsiz is the same as described above.

Rmax is the maximum size of any record in the file.

Rnum is the number of records that you want in each bucket. Variable-length records in an indexed file bucket always occupy Rmax plus nine bytes of record control information. Fifteen bytes are required for bucket control information.

SPECIAL NOTE FOR INDEXED FILES

If the BKS field is not specified and a maximum record size (MRS) is specified, then VAX-11 RMS will use a bucket size to ensure that at least one maximum size record will fit. Generally, performance on record insertion and sequential retrieval on primary key is improved if at least 3 or 4 data records will fit into a primary data bucket. If either the bucket size or the disk cluster size is other than 1 block, then it is advisable to use a default extend quantity (DEQ) which is the least common multiple of the bucket size and cluster size, to avoid unused, but allocated blocks within the file.

\$FAB BLS

4.2.4 Block Size

The BLS parameter is used as input only for magnetic tape files. When you create a magnetic tape file, you can set the block size field before you issue the \$CREATE macro instruction. In all other cases, VAX-11 RMS ignores it. When you open an existing file with a \$OPEN macro instruction, VAX-11 RMS returns the device buffer size if the file is organized sequentially. For terminals, this is the value of the WIDTH setting. For mailboxes, this is the value of the maximum message size.

Format

BLS=block-size

block-size

The size, in bytes, of the blocks on the tape, in the range of 20 through 65532. If this parameter is 0, the default selected when the volume was mounted is used.

For compatibility with RMS-11, block size is always rounded off to be a multiple of 4. For example, if you set the block length to 38, you would get 40.

For example, to set the block length to 4096, the syntax is:

```
$FAB      BLS=4096,MRS=132
```

NOTE

To create a magnetic tape for interchange with other DIGITAL operating systems (non-VAX/VMS), you should consult the documentation for the target system regarding possible limitations on block size. To ensure compatibility with non-DIGITAL systems, ANSI standards require that the block size be less than or equal to 2048 bytes.

\$FAB CTX

4.2.5 User Context

The CTX parameter conveys user information to a completion routine in your program. The user context field set by this parameter is intended solely for your use; VAX-11 RMS never uses it for record management activities.

THE FILE ACCESS BLOCK

Format

CTX=value

value

represents any user-specified value, up to four bytes long.

For example, to pass along the symbolic value T1DONE, the syntax is:

```
$FAB      CTX=T1DONE,BKS=4
```

\$FAB DEQ

4.2.6 Default File Extension Quantity

The DEQ parameter sets the default file extension quantity field, which specifies the number of blocks to add when a disk file is extended automatically. This automatic extension occurs whenever your program performs an operation with a \$PUT or \$WRITE macro instruction and the currently allocated space is exhausted.

Format

DEQ = extension-quantity

extension-quantity

The number of blocks to be added when automatic extension is required. This number must be in the range of 0 through 65,535 and is rounded up to the next cluster boundary. If you specify 0, the file will be extended using a VAX-11 RMS determined default extension value.

For example, to specify a default extension quantity of 80 blocks, the syntax is:

```
$FAB DEQ=80
```

User Considerations

1. When creating a new file, you can specify the extension quantity for the file by setting the desired value in the default extension quantity field before issuing a \$CREATE macro instruction. This value becomes a permanent attribute for the file.
2. When processing an existing file, you can temporarily override the default extension quantity specified when the file was created. To do this, set the desired value before issuing the \$OPEN macro instruction. Once the file is closed, the default extension quantity reverts to the value set when the file was created.
3. See notes under BKS for indexed files.

NOTE

The use of an allocation XAB will override the value in this field. See chapter 6 for a detailed description of allocation XAB's.

\$FAB DNA

4.2.7 Default File Specification String Address

You can use the DNA parameter to set program defaults in the default file specification string address field for the missing components (if any) of the file specification string pointed to by the file specification string address field. This parameter works with the DNS parameter, which initializes the default file specification string size (see Section 4.2.8).

The default file specification string is used primarily when accepting file specifications interactively; file specifications known to a user program are normally completely specified in the file specification string address and size fields (the FNA and FNS parameters). You can specify defaults for one or more of the following file specification components:

- Node
- Device
- Directory
- File name
- File type
- File version number

Format

DNA = address

address

The symbolic address of an ASCII string containing one or more components of a file specification. The components in the string must be in the order in which they would occur in a complete file specification.

For example, assume an ASCII string is stored at a memory location whose symbolic address is DFNAM. To store the address of this string in the default file specification string address field, so that DFNAM will be used during execution of a \$OPEN or \$CREATE macro instruction, the syntax is:

```
$FAB DNA=DFNAM,DNS=4
```

This default file specification string address is only effective if the components are missing from the string whose address is stored in the file specification address field.

Section 4.2.9 describes a simpler technique -- using the DNM parameter -- for setting the default file specification string address at assembly time.

\$FAB DNS

4.2.8 Default File Specification String Size

The DNS parameter sets a value in the default file specification string size field. This value indicates the size, in bytes, of the string whose address is contained in the default file specification string address field.

THE FILE ACCESS BLOCK

Format

DNS=value

value

A symbolic or numeric value representing the size of the default file specification string. The numeric value is in the range of 1 through 255.

For example, assume that your program contains the directive:

```
DFNAM: .ASCII /.DAT/
```

The following DNS parameter would set the default file specification string size field:

```
$FAB DNS=4,DNA=DFNAM
```

Section 4.2.9 describes another technique -- using the DNM parameter -- for setting the default file specification string size.

\$FAB DNM

4.2.9 Default File Specification

The DNM parameter sets two fields in the FAB: the default file specification string address (DNA) and the default file specification string size (DNS). The specified default file specification string is stored in the special program section \$RMSNAM.

Format

```
DNM=<FILESPEC>
```

<FILESPEC>

The ASCII default file specification string. The angle brackets (<>) are required syntax.

For example:

```
$FAB DNM=< .DAT >
```

\$FAB FAC

4.2.10 File Access

The FAC parameter initializes the file access field. You must indicate to VAX-11 RMS what types of operations you intend to perform on the file. After you open a file, VAX-11 RMS rejects any operation your program attempts if that operation was not specified in the file access field when you issued a \$OPEN or \$CREATE macro instruction for the file.

THE FILE ACCESS BLOCK

If your program will issue any of the following macro instructions, you must specify them by setting the file access field for the appropriate operation:

- \$DELETE
- \$FIND
- \$GET
- \$PUT
- \$READ
- \$SPACE
- \$TRUNCATE
- \$UPDATE
- \$WRITE

Format

FAC=<BIO,BRO,DEL,GET,PUT,TRN,UPD>

BIO

Used for block I/O operations involving a \$READ or \$WRITE macro instruction, with Get and Put access, respectively, and also with a \$SPACE macro instruction. Furthermore, specifying block I/O prohibits the use of any record I/O operations (GET, PUT, DEL, UPD, TRN).

BRO

Similar to BIO, except that record I/O operations are also allowed.

DEL

Allows operations with a \$DELETE macro instruction.

GET

Allows operations with a \$GET or \$FIND macro instruction. This is the default when you are opening this file and either the FAC parameter is not specified or the DEL, UPD, or TRN operations are specified on the FAC parameter. If you specify GET with either BIO or BRO, you can perform operations with a \$READ macro instruction.

PUT

Allows operations with a \$PUT macro instruction. This will be the default if you are creating this file. If you specify PUT with either BIO or BRO, you can perform operations with a \$WRITE macro instruction.

TRN

Allows operations with a \$TRUNCATE macro instruction. Also allows use of the truncate put (TPT) record option on a \$PUT and \$WRITE macro instruction (see Section 5.2.14).

THE FILE ACCESS BLOCK

UPD

Allows operations with a \$UPDATE macro instruction. This is also required for Put operations with the update if (UIF) option on indexed files.

You may specify more than one operation with the FAC parameter. However, if you do, the group of operations must be enclosed in angle brackets; when only one operation is specified, no angle brackets are needed. Multiple operations can be specified in any order. For example, <GET,PUT,UPD> or <UPD GET PUT>.

The following example indicates that operations with a \$PUT macro instruction are going to be performed.

```
$FAB    FAC=PUT,ALQ=132,DEQ=16
```

A request for operations with \$GET, \$PUT, and \$UPDATE macro instructions would be specified as follows:

```
$FAB    FAC=<GET,UPD,PUT>
```

Each operation has its own symbolic bit offset and mask value.

\$FAB FNA

4.2.11 File Specification String Address

The FNA parameter initializes the file specification string address field. This parameter works with the FNS parameter, which initializes the file specification string size field (see Section 4.2.12). The file specification string address contains the address of an ASCII string that specifies the path to a file to be processed. If this string does not contain all the components of a full file specification, VAX-11 RMS will use the defaults supplied in the default file specification string (see Sections 4.2.7, 4.2.8, and 4.2.9). If no default string is present, or if the file specification is still incomplete, VAX-11 RMS provides further defaulting (see Section 8.2).

Format

```
FNA=address
```

address

The symbolic address of an ASCII string containing the file specification.

For example, assume that the following directive is in your program:

```
FLNAM:  .ASCII  /MASTER.OLD/
```

The syntax for the FNA parameter is:

```
$FAB    FNA=FLNAM,FNS=10
```

See 4.2.13 for an alternate method of setting the file specification at assembly time with the FNM macro.

\$FAB FNS

4.2.12 File Specification String Size

The FNS parameter initializes the file specification string size field. This field describes the length, in bytes, of the ASCII string pointed to by the file specification string address field (FNA).

Format

FNS=value

value

A numeric or symbolic value representing the size, in bytes, of the file specification string, in the range of 0 through 255.

For example, assume that the following directive is in your program:

FLNAM: .ASCII /INPUTFILE:/

The syntax for the FNS parameter is:

\$FAB FNS=10,FNA=FLNAM

Section 4.2.13 describes another technique -- using the FNM parameter -- for setting the file specification string size field.

\$FAB FNM

4.2.13 File Specification

The FNM parameter sets two fields in the FAB: the file specification string address and the file specification string size. It causes the specified string to be stored in the special program section named \$RMSNAM.

Format

FNM = <FILESPEC>

<FILESPEC>

The ASCII file specification string; the angle brackets (<>) are required syntax.

For example:

\$FAB FNM=DISK:[DATA]<FILE.DAT>,ALQ=132

\$FAB FOP

4.2.14 File Process Options

The FOP parameter sets indicators in the file-processing options field that represent requests for optional file-handling operations.

THE FILE ACCESS BLOCK

Format

FOP=<CBT,CIF,CTG,DFW,DLT,MXV,NAM,NEF,NFS,OFF,POS,
RCK,RWC,RWO,SCF,SPL,SQO,SUP,TEF,TMD,TMP,UFO,WCK>

With the exception of the CBT, CTG, RCK, and WCK bits, the contents of this field are not modified by VAX-11 RMS operations.

Each option is interpreted as follows:

Allocation and Extension Options:

CBT

Contiguous best try: indicates that the file is to be allocated contiguously on a "best effort" basis. It is input to the create service, and is output from the open service to indicate the file status. Note that the file will take on the contiguous best try attribute only if a space allocation is actually performed. The CBT option takes precedence over the CTG option (below).

CTG

Contiguous; indicates that the space for the file is to be allocated contiguously. If this cannot be done, the operation fails. It is input to the create service, and is output by the open service to indicate the status of the file. The CBT option (above) takes precedence over the CTG option.

TEF

Truncate at end of file; indicates that unused space allocated to a file is to be deallocated on a close service. This applies to sequential files only.

Performance Options:

SQO

Sequential only; indicates that this file can be processed sequentially only, thus allowing certain processing optimizations. Any attempt to perform random access will result in an error. This option is input to the create and open services and applies to the Find, Get, and Put services for sequential files. At present, this performance enhancement applies only to network operations.

DFW

Deferred write; indicates that writing back to the file of modified I/O buffers is to be deferred until the buffer must be used for other purposes. This option applies to relative files and indexed files.

Reliability Options:

RCK

Read-check; specifies that transfers from disk volumes are to be checked by a follow-up read-compare operation. This is an input to the open and create services. If RCK is set, then checking is performed for the duration of the access. The RCK is also an output of the open service, which indicates the default for the file.

THE FILE ACCESS BLOCK

WCK

Write-check; indicates that transfers to disk are to be checked by a follow-up read-compare. Similar to the RCK option.

Filename Parsing Modifiers:

CIF

Create if; causes the file to be opened if it already exists. If the file does not exist, it is created and the alternate success RMS\$ CREATED is returned. It is input only on a create service. The CIF option takes precedence over the SUP option.

MXV

Maximize version; indicates that the version number of the file should be the maximum of the explicit version number given in the file specification or one greater than the highest version number for an existing file in the same directory with the same file name and file type.

NAM

NAM block inputs; indicates that the NAM block specified in the name block address field is to be used to provide:

- The device identification, file identification, and/or the directory identification when the file is being opened, closed, or deleted
- The device identification and the directory identification when the file is being created

For further detail, see Section 8.2.

OFF

Output file parse; specifies that the related file resultant file specification string, if used, is to provide file name and file type defaults only (see Section 8.2).

SUP

Supersede; allows an existing file to be superseded on a create service by a new file of the same name, type, and version. The CIF option (above) takes precedence over the SUP option.

File Disposition Options:

DLT

Delete; indicates that the file is to be deleted when it is closed; this option may be specified on a close, create, or open service. You can specify the DLT option with the SCF or SPL option. However, if you do not have a NAM block in conjunction with this, the file's directory entry will not be removed.

SCF

Submit command file; indicates that the file is to be submitted as a batch-command file to the process-default batch queue when the file is closed. This option can be specified for the close, create, or open services. It is currently implemented for sequential files only.

THE FILE ACCESS BLOCK

SPL

Spool; indicates that the file is to be spooled to the process default print queue when the file is closed. When using this option, you should normally use a NAM block and specify the NAM option (of this file-processing options field) so that the resultant file specification string is available. This option can be specified for the Close, Create, or Open services. It is currently implemented for sequential files only.

TMD

Temporary marked for delete; indicates that a temporary file is to be created, and then deleted when the file is closed. This option is input only to the create service. The TMD option takes precedence over the TMP option (below).

TMP

Temporary; indicates that a temporary file is to be created and retained, but that no directory entry will be made for this file. This option is input only to the create service. The TMD option (above) takes precedence over the TMP option.

Magnetic Tape Processing Options:

NEF

Not end of file; inhibits the positioning to the end of file when a tape file is opened and the file access field of this FAB indicates a Put operation.

POS

Current position; indicates that the magnetic tape volume set should be positioned immediately after the most recently closed file when the next file is created. However, if the RWO option of this field is also set, it overrides the POS option and positions to the beginning of the volume set.

RWC

Rewind on close; specifies that the magnetic tape volume is to be rewound when the file is closed. This option can be specified for the Close, Create, or Open services.

RWO

Rewind on open; specifies that the magnetic tape volume is to be rewound before the file is opened or created. The RWO option takes precedence over the POS option (above).

Non-Standard Processing Options:

NFS

Non-file-structured; indicates on Open or Create that the volume is to be processed in a non-file-structured manner. This allows the use of volumes created on non-DIGITAL systems. For further explanation, see Section 12.4.

THE FILE ACCESS BLOCK

UFO

User file open; indicates that VAX-11 RMS will open or create the file only. No further VAX-11 RMS operations can be done with this file. To perform any further processing on the file, you must use the QIO system service with the channel number that is returned in the status value field (STV). This channel will be assigned in the mode of the caller. For the create service, the end of file mark will be set to the end of the block specified in the allocation options field on input (see Section 4.2.2). For either the open or create services, the IFI field is set to 0 on return to indicate that VAX-11 RMS cannot perform any more operations on the file. If you use the UFO on \$OPEN or \$CREATE, the channel needs only to be deassigned when you are finished with the file. A Close operation is not required.

You can specify more than one option with the FOP parameter. However, if you do, you must enclose the group of options in angle brackets. When you specify only one option, no angle brackets are needed. The options can be specified in any order.

For example, to rewind a tape file as part of the close operation, the syntax is:

```
$FAB    BLS=4096,FOP=RWC
```

Each option has its own symbolic bit offset and mask value.

\$FAB FSZ

4.2.15 Fixed Control Area Size

The FSZ parameter initializes the fixed control area size field, which is used when dealing with variable with fixed-length control records. When you create a file with this type of record, you must set the value for the fixed-control area before you issue the \$CREATE macro instruction. When you open an existing file that contains variable with fixed control records, VAX-11 RMS sets this field equal to the value specified when the file was created. The FSZ parameter is not applicable to indexed files.

Format

```
FSZ=header-size
```

header-size

The numeric value, in bytes, of the size of the fixed control area, in the range of 1 to 255. The default size is 2 bytes. If you specify 0, then the default size is used.

For example, if each variable with fixed-length control record is to have an 8-byte fixed control area, the syntax is:

```
$FAB    FOP=WCK,FSZ=8
```

\$FAB MRN**4.2.16 Maximum Record Number**

The MRN parameter sets the maximum record number field, which indicates the highest record number that can be written into this file. You can use this parameter only for relative files. If you attempt to put or get a record with a higher relative record number than the specified limit, an error will occur and VAX-11 RMS will return a message indicating an invalid record number. If, however, you specify 0, checking is suppressed.

Format

MRN=max-rec-number

max-rec-number

Numeric value of the highest numbered record allowed in the file, in the range of 0 to 2,147,483,647. The default for this parameter is 0.

For example, to set the highest relative record number to 10000, the syntax is:

\$FAB MRN=10000,FOP=WCK

NOTE

VAX-11 RMS does not maintain the relative record number of the highest existing record in the file.

\$FAB MRS**4.2.17 Maximum Record Size**

The MRS parameter sets the maximum record size field, which indicates, in bytes, the size of the records in the file.

For fixed-length records, the value represents the actual size of each record in the file. You must specify a size when you create a file with fixed-length records.

For variable-length records, the value represents the size of the largest record that can be written into the file. If the file is not a relative file, a value of 0 is used to suppress record size checking, thus indicating that there is no user limit on record size. However, the record size must conform to physical limitations. In the case of indexed and relative files, for example, records may not cross bucket boundaries.

For variable with fixed-length control records, the value includes only the data portion; it does not include the size of the fixed control area.

For all relative files, the size is used to determine the size of the record cell, and is used in conjunction with the bucket size field (see Section 4.2.3).

THE FILE ACCESS BLOCK

You specify a value when you issue a \$CREATE macro instruction. VAX-11 RMS returns the maximum record size when you issue a \$OPEN macro instruction.

Format

MRS=max-rec-size

max-rec-size

The following table summarizes the maximum record size allowed for the various file and record formats:

FILE ORGANIZATION	RECORD FORMAT	MAXIMUM ALLOWED
Sequential	Fixed-length	32,767
Sequential (Disk)	Variable-length	32,767-FSZ ¹
Sequential (ANSI Tape)	Variable-length	9,995-FSZ
Relative	Fixed-length	16,383
Relative	Variable-length	16,381-FSZ
Indexed Sequential	Fixed-length	16,362
Indexed Sequential	Variable-length	16,360

For example, to set a maximum record size of 512 bytes, the syntax is:

```
$FAB MRS=512,MRN=10000
```

NOTE

The length of the largest record actually existing in a sequential file with variable or VFC record format is also maintained by VAX-11 RMS and is available through the file header characteristics XAB (LRL field of \$XABFHC) (see Section 6.9).

\$FAB NAM

4.2.18 Name Block Address

The NAM parameter lets you set a symbolic address in the name block address field of the FAB. This address points to the NAM block you want to use when performing an operation, such as an open or create, on a file. The NAM block, described in Chapter 7, is required only in conjunction with the file specification processing macro instructions (see Chapter 13).

Format

NAM=nam-address

nam-address

The symbolic address of the NAM block.

1. The FSZ represents the size of the fixed control area of a record. The FSZ=0 for variable-length records. The FSZ is equal to the size, in bytes, for the fixed control area of the VFC (variable with fixed-length control) records.

THE FILE ACCESS BLOCK

For example, if a \$NAM macro instruction for a NAM block has a label of NMBLK, the syntax is:

```
$FAB MRS=512,MRN=1000,NAM=NMBLK
```

\$FAB ORG

4.2.19 File Organization

The ORG parameter sets the file organization field, indicating the arrangement of the data in the file. You must set this field before you issue a \$CREATE macro instruction. VAX-11 RMS returns the contents of this field when you issue a \$OPEN macro instruction.

Format

$$\text{ORG} = \left\{ \begin{array}{l} \text{REL} \\ \text{IDX} \\ \text{SEQ} \end{array} \right\}$$

REL

Relative file organization.

IDX

Indexed file organization.

SEQ

Sequential file organization. This is the default.

For example, to set the file organization field to relative, the syntax is:

```
$FAB MRN=1000,ORG=REL,MRS=512
```

Each organization has its own symbolic value.

- REL -- FAB\$C_REL
- IDX -- FAB\$C_IDX
- SEQ -- FAB\$C_SEQ

\$FAB RAT

4.2.20 Record Attributes

The RAT parameter initializes the record attributes field with special control information pertaining to the records in the file. If you need this information, set this field before you issue a \$CREATE macro instruction. VAX-11 RMS sets the field when you issue a \$OPEN macro instruction.

Format

$$\text{RAT} = \langle \text{BLK} \left\{ \begin{array}{l} \text{CR} \\ \text{FTN} \\ \text{PRN} \end{array} \right\} \rangle$$

THE FILE ACCESS BLOCK

BLK

Indicates that records do not cross block boundaries. This information applies to sequential files only.

CR

Indicates that each record is to be preceded by a line feed and followed by a carriage return when the record is written to a carriage control device such as a line printer or terminal.

FTN

Indicates that the first byte of each record contains a FORTRAN (ASA) carriage control character, defined as follows:

Byte 0 Value (hexadecimal)	ASCII Character	Meaning
0	(null)	Null carriage control. (Sequence: print buffer contents.)
20	(space)	Single-space carriage control. (Sequence: newline, print buffer contents, RETURN.)
30	0	Double-space carriage control. (Sequence: newline, newline, print buffer contents, RETURN.)
31	1	Page eject carriage control. (Sequence: form feed, print buffer contents, RETURN.)
28	+	Overprint carriage control. (Sequence: print buffer contents, RETURN.) Allows double printing for emphasis.
24	\$	Prompt carriage control. (Sequence: newline, print buffer contents.)
All other values		Same as ASCII space character: single-space carriage control.

PRN

Indicates the print file format for variable with fixed-length control records, where the fixed control area contains the print file information, including carriage control. The first byte of the fixed control area constitutes a "prefix" area, and the second byte constitutes a "postfix" area, specifying carriage control to be performed before and after printing the record respectively. The encoding scheme of both bytes is as follows (even though they are interpreted separately):

Bit 7	Bits 0-6	Meaning
0	0	No carriage control is specified, that is, NULL.
0	1-7F	Bits 0 through 6 are a count of newlines (line feeds followed by carriage return).

THE FILE ACCESS BLOCK

Bit 7	Bit 6	Bit 5	Bits 0-4	Meaning
1	0	0	0-1F	Output the single ASCII control character specified by the configuration of bits 0 through 4 (7-bit character set).
1	1	0	0-1F	Output the single ASCII control character specified by the configuration of bits 0 through 4 which are translated as ASCII characters 128 through 159 (8-bit character set).
1	1	1	0-1F	Reserved

Only the BLK attribute can be paired with another attribute. You cannot use CR, FTN, and PRN together in any combination. When BLK is used with another attribute, you can specify them in any order; the angle brackets are part of the required syntax when BLK is used with another attribute.

The following example indicates that records do not cross block boundaries.

```
FAB$    ORG=SEQ,RAT=BLK
```

Each option has its own symbolic bit offset and mask value.

\$FAB RFM

4.2.21 Record Format

The RFM parameter initializes the record format field to indicate the type of records in the file. When you create the file, you must set this field before you issue the \$CREATE macro instruction. VAX-11 RMS returns the record format when you issue a \$OPEN macro instruction.

Format

$$\text{RFM} = \left\{ \begin{array}{l} \text{FIX} \\ \text{VAR} \\ \text{VFC} \\ \text{UDF} \end{array} \right\}$$

FIX

Indicates fixed-length record format.

VFC

Indicates variable-length with fixed-length control record format. This format is not valid for indexed files.

VAR

Indicates variable-length record format. This is also the default value (assembly time default).

THE FILE ACCESS BLOCK

UDF

Indicates undefined record format. The undefined record format is valid for sequential file organization only, and can be processed only through the use of block I/O. This is the default value if the FAB is not initialized with a \$FAB macro instruction.

For example, to indicate that records are fixed-length, the syntax is:

```
$FAB    RFM=FIX,FAC=GET
```

Each record format has its own symbolic value.

- FIX -- FAB\$C_FIX
- VAR -- FAB\$C_VAR
- VFC -- FAB\$C_VFC
- UDF -- FAB\$C_UDF

\$FAB RTV

4.2.22 Retrieval Window Size

The RTV parameter initializes the retrieval window size field. This field identifies the number of retrieval pointers you want VAX-11 RMS to maintain in memory for the file.

Format

```
RTV=window-size
```

window-size

The number of retrieval pointers, in the range of 0 through 127, or 255. A value of 0 indicates that VAX-11 RMS is to use the system default number of retrieval pointers. A value of 255 means to map the entire file, if possible. Values between 128 and 254 inclusive are reserved for future use.

For example, to reserve ten retrieval pointers, the syntax is:

```
$FAB    FAC=GET,RTV=10,RFM=FIX
```

\$FAB SHR

4.2.23 File Sharing

The SHR parameter sets a value in the file-sharing field, indicating the operations other users can perform when they are sharing access to the file with you. VAX-11 RMS supports file sharing for all relative and indexed file operations, as well as for sequential files with 512-byte fixed-length records. For additional information concerning file sharing, see Chapter 10.

THE FILE ACCESS BLOCK

Format

SHR=<PUT,GET,DEL,UPD,NIL,UPI,MSE>

PUT

Allows other users to write records to the file.

GET

Allows other users to read the file.

DEL

Allows other users to delete records from the file.

UPD

Allows other users to update records that currently exist in the file.

NIL

Prohibits any type of file sharing by other users. (If specified along with other operations, NIL takes precedence.)

UPI

Allows one or more writers for a sequential file or a shared file which is open for block I/O. The user assumes the responsibility for any required interlocking. This operation is set in combination with PUT, GET, UPD, and/or DEL, but does not apply to relative and indexed files.

MSE

Allows multistream access. You must specify MSE whenever you are going to issue \$CONNECT macro instructions for multiple RABs for this FAB. This option is not available for sequential files with other than 512-byte fixed-length records.

You can specify one or more file-sharing operations in any order.

For example, to allow read, write, and delete operations by other users, the syntax is:

```
$FAB    RTV=10,RFM=FIX,SHR=<DEL,PUT,GET>
```

Each file-sharing operation has its own symbolic bit offset and mask value.

THE FILE ACCESS BLOCK

NOTE

If you do not specify the SHR, VAX-11 RMS enters a value of 0 in the file-sharing field. Defaults apply as follows:

- If the file access field (FAC parameter) is set or defaulted to GET, the file-sharing field is defaulted to GET.
- If the file access field is set or defaulted to either PUT, DEL, UPD, or TRN, the file-sharing field is defaulted to NIL.

\$FAB XAB

4.2.24 Extended Attribute Block Pointer

For some operations, you must associate Extended Attribute Blocks (XABs) with a FAB to convey additional attributes about a file (see Chapter 6 for a description of an XAB). The XAB parameter sets the extended attribute block pointer field with the address of the first associated block (of a potential chained list of such blocks) for the file.

Format

XAB=xab-address

xab-address

The symbolic address of the first XAB. A value of 0 (the default) indicates no XABs for the file.

For example, if the \$XAB macro instruction has a label of HDRXAB, the syntax is:

```
$FAB XAB=HDRXAB
```

NOTES

1. If you specify an XAB for either a \$OPEN or \$DISPLAY macro instruction, VAX-11 RMS returns the attributes for the file to the XAB.
2. If you specify an XAB for a \$CLOSE, \$CREATE, or \$EXTEND macro instruction, VAX-11 RMS uses the XAB as input to those functions.

THE FILE ACCESS BLOCK

4.3 NONINITIALIZABLE FAB FIELDS

The following list describes the FAB fields that you cannot initialize. Either they are statically initialized, or VAX-11 RMS sets them for you.

Output Only Fields:

DEV

The device characteristics field is set by VAX-11 RMS when you issue an \$OPEN, \$CREATE or \$PARSE macro instruction. This field allows VAX-11 RMS to communicate to your program the generic characteristics of the device containing the file. Although you cannot initialize this field at assembly time, you can interrogate the contents of the fields through the symbolic offsets. Table 4-2 lists the bits in the device characteristics field. Each bit described in this table has its own symbolic bit offset and mask value. These definitions can be made available to your program by issuing the \$DEVDEF macro instruction. The bit offset is formed by prefixing the characteristic name with DEV\$V_. For example:

```
REC -- DEV$V_REC
```

The mask value is formed by prefixing the characteristic name with DEV\$M_. For example:

```
REC -- DEV$M_REC
```

SDC

Spooling device characteristics field; equivalent to the device characteristics field (DEV), except that spooling device characteristics refer to the intermediate device used for spooling. The bit definitions for the SDC field are the same as those defined for the DEV field.

STS

Completion status code field; VAX-11 RMS sets this field with success or failure codes before control is returned to your program. Register 0 will contain the same status. Potential error codes for specific operations are listed under their descriptions in the chapters on file and file specification processing (9 and 13). A complete list of all RMS error codes is in Appendix A. Status codes are further discussed in Section 8.4.

STV

Status value field; communicates additional completion information to your program, based on the type of operation performed and the contents of the completion status code field. See Appendix A for the instances when VAX-11 RMS uses the status value field. For more information on completion codes, see Section 8.4.

THE FILE ACCESS BLOCK

Internal File Identifier:

IFI

Internal file identifier field; associates the FAB with the corresponding internal file access block. It is set by VAX-11 RMS on successful Create or Open services. It is then an input for subsequent Close, Display, and Extend operations. The Close service deallocates the internal control structures and clears the IFI. When the user file open (UFO) option in the file options (FOP) field is specified, no internal structures are allocated on the Create or Open service. Therefore, the IFI will remain cleared.

Static Fields:

BID

Block identifier field; identifies the block as a FAB to VAX-11 RMS. This field is set by the \$FAB macro instruction to the symbolic value FAB\$C_BID, and must not be altered.

BLN

Block length field; defines the length of the FAB to VAX-11 RMS. This field is set by the \$FAB macro instruction to the symbolic value FAB\$C_BLN, and must not be altered.

THE FILE ACCESS BLOCK

Table 4-2
Device Characteristics

Bit Name	Description
ALL	Device is allocated
AVL	Device is available for use
CCL	Carriage control device
DIR	Directory structured device
DMT	Device is marked for dismount
ELG	Device is error log enabled
FOD	File-oriented device (disk and magnetic tape)
FOR	Device is mounted foreign (i.e., non-file structured)
GEN	Device is a generic device
IDV	Device can provide input
MBX	Device is mailbox
MNT	Device is currently mounted
NET	Network device
ODV	Device can accept output
RCK	Device has read check enabled
REC	Record-oriented device (terminal, line printer, etc.). If field is 0, device is assumed to be block-oriented (disk, magnetic tape). All record-oriented devices are considered sequential in nature.
RND	Device is random access in nature
RTM	Device is realtime in nature; not suitable for VAX-11 RMS usage
SDI	Single directory device (master file directory only)
SHR	Shareable device
SPL	Device is being spooled
SQD	Sequential block-oriented device (magnetic tape)
SWL	Device is currently software write-locked
TRM	Terminal device
WCK	Device has write check enabled

CHAPTER 5

THE RECORD ACCESS BLOCK

This chapter describes the Record Access Block (RAB), the fields in the RAB, and the parameters of the \$RAB macro instruction. All record operations (described in Chapter 11) and Block I/O operations (described in Chapter 12) require the RAB as a parameter.

5.1 THE PURPOSE OF THE RECORD ACCESS BLOCK

The RAB is the second type of user control block that you allocate, either at assembly time or run time, to communicate with VAX-11 RMS. During program execution, you associate a RAB with a File Access Block (FAB) to establish a record stream using a \$CONNECT macro instruction. Once you have established a record stream, you use the fields of the RAB to define to VAX-11 RMS the next record you want to access in the file.

Each RAB is linked to a FAB, and represents a record request stream on the file associated with the FAB. Once you establish this link, you can use the fields of the RAB to define for VAX-11 RMS the next logical record you want to access and various characteristics about that record.

You allocate a RAB with a \$RAB macro instruction, and initialize the fields either at assembly time (through keyword parameters) or by direct manipulation at run time. You initialize the RAB at run time through either keyword parameters with the \$RAB STORE macro instruction (see Chapter 14) or the defined symbolic offsets. You need one RAB for each record stream in your program.

Each field in the RAB has a 3-character mnemonic name. All access to these fields is through this name (by keyword or offset). However, some of the fields, as in the FAB, are static or output only; therefore, you need not initialize them. Table 5-1 summarizes the fields of the RAB, including the static and output-only fields.

THE RECORD ACCESS BLOCK

Table 5-1
Record Access Block Fields

Field & Keyword Name	Field Size	Description	Offset
BID ²	byte	Block identifier	RAB\$B_BID
BKT	longword	Bucket code	RAB\$L_BKT
BLN ²	byte	Block length	RAB\$B_BLN
CTX	longword	Context	RAB\$L_CTX
FAB	longword	File access block address	RAB\$L_FAB
ISI ¹	word	Internal stream identifier	RAB\$W_ISI
KBF	longword	Key buffer address	RAB\$L_KBF
KRF	byte	Key of reference	RAB\$B_KRF
KSZ	byte	Key size	RAB\$B_KSZ
MBC	byte	Multiblock count	RAB\$B_MBC
MBF	byte	Multibuffer count	RAB\$B_MBF
PBF	longword	Prompt buffer address	RAB\$L_PBF
PSZ	byte	Prompt buffer size	RAB\$B_PSZ
RAC	byte	Record access mode	RAB\$B_RAC
RBF	longword	Record address	RAB\$L_RBF
RFA ¹	3 words	Record's file address	RAB\$W_RFA
RHB	longword	Record header buffer	RAB\$L_RHB
ROP	longword	Record-processing options	RAB\$L_ROP
RSZ	word	Record size	RAB\$W_RSZ
STS ¹	longword	Completion status code	RAB\$L_STS
STV ¹	longword	Status value	RAB\$L_STV
STV0 ³	word	low-order word status value	RAB\$W_STV0
STV2 ³	word	high-order word status value	RAB\$W_STV2
TMO	byte	Timeout period	RAB\$B_TMO
UBF	longword	User record area address	RAB\$L_UBF
USZ	word	User record area size	RAB\$W_USZ

¹ Indicates nonuser-initialized field.

² Indicates statically initialized field (by the \$RAB macro instruction) to identify this control block as a RAB.

³ Alternate definition of STV field.

\$RAB

5.2 RAB ALLOCATION

The format of the \$RAB macro instruction is shown below. Every parameter is optional, depending on the function to be performed and the combination of parameters in the macro instruction as a whole.

Format:

OPERATION	PARAMETERS
label: \$RAB	BKT=number CTX=value FAB=fab-address KBF=buffer-address KRF=key-number KSZ=size MBC=blocks MBF=buffers PBF=prompt-address PSZ=prompt-size RAC= { SEQ } { KEY } { RFA } RBF=buffer-address RHB=header-address ROP=<ASY BIO CCO CVT EOF KGE KGT LIM LOA LOC NLK NXR PMT PTA RAH RLK RNE RNF TMO TPT UIF ULK WBH> RSZ=record-size TMO=seconds UBF=buffer-address USZ=buffer-size

The \$RAB macro instruction allocates and initializes storage for a RAB. You cannot use this macro instruction within a sequence of executable instructions. In some cases, specific default values are assigned automatically when you omit a parameter. These specific defaults are noted in the text that explains each parameter. If there is no specific default, VAX-11 RMS uses a default value of 0.

label: \$RAB

5.2.1 Label

The label for the \$RAB macro instruction lets you name a RAB, and thereby provides symbolic access to a particular RAB within your program. The label is optional but, when used, must precede the symbol \$RAB and be separated from \$RAB by a colon (:). For example:

```
INPUT: $RAB
```

\$RAB BKT

5.2.2 Bucket Code

The BKT parameter initializes the bucket code field of the RAB. This field is used as follows:

1. With records in a relative file
2. When performing block I/O

For relative files, the relative record number of the record acted upon (or which produced an error) is returned to the bucket code field only after the completion of a sequential operation. That is, VAX-11 RMS returns the relative record number when you set the record access mode for sequential access (RAC=SEQ) on the execution of a \$GET, \$PUT, or \$FIND macro instruction.

When performing block I/O on disk devices, you must store (in the bucket code field) the virtual block number (VBN) of the first block you want to read or write. For all other devices, this field is not used. If you specify a VBN of 0, VAX-11 RMS will begin the block transfer at the block pointed to by the Next Block Pointer (NBP). The NBP is an internal pointer maintained by VAX-11 RMS, and is described in Chapter 12.

Format

```
BKT=number
```

number

A relative record number or a numeric value representing the virtual block number to be accessed.

For example, to indicate access to the tenth block of the file when the program performs its first block I/O operation, the syntax is:

```
$RAB BKT=10,CTX=RECOK
```

\$RAB CTX**5.2.3 Context**

The CTX parameter initializes the context field, which is a field devoted exclusively to your use. VAX-11 RMS makes no use of the contents of this field; therefore, you can set any value you want in this field. For example, you could use this field to communicate with a completion routine in your program.

Format

CTX=value

value

Any user-selected value, up to one longword in length.

For example, to initialize the context field to the value of the symbol RECOK, the syntax is:

\$RAB CTX=RECOK,BKT=10

\$RAB FAB**5.2.4 File Access Block Address**

The FAB parameter initializes the file access block address field of the RAB. When you issue a \$CONNECT macro instruction, you must set this field to indicate the address of the FAB associated with the open file.

Format

FAB=fab-address

fab-address

The symbolic address of the FAB for the file.

For example, if you define the label of the FAB for the file as MASTER, the syntax is:

\$RAB FAB=MASTER,CTX=RECOK

\$RAB KBF**5.2.5 Key Buffer Address**

The KBF parameter initializes the key buffer address field. You use this field when the record access mode (RAC) field specifies random access by key value (see Section 5.2.12), and you set it to the address of the buffer that contains the key of the desired record. For a relative file (or for a sequential disk file with fixed-length records), the key is the relative record number. For an indexed file, the key is the key value within the record for the key of reference (KRF) (see Section 5.2.6).

THE RECORD ACCESS BLOCK

Format

KBF=buffer-address

buffer-address

The symbolic address of the buffer containing the key.

For example, if the label of the buffer that provides the relative record number is RELKEY, you initialize the KBF parameter as follows:

```
$RAB KBF=RELKEY,CTX=RECOK
```

NOTE

Before issuing a \$GET or \$FIND macro instruction in random mode to an indexed file, you place in KBF the address of a location containing a key value. The size of this key value must be specified in the KSZ field. During execution of the Get or Find operation, VAX-11 RMS uses the key value described by the KBF and KSZ fields to search an index (which you specify through the contents of the KRF field of the RAB) and locate the desired record in the file. The type of match (that is, exact, generic, approximate, or approximate and generic) that VAX-11 RMS attempts between the key value you specify and key values in records of the file is determined by the KSZ field and the ROP field.

NOTE

The key buffer address field uses the same location in the RAB as the prompt buffer address field. There is no conflict between these two fields, however, because the prompt buffer address field is used only for terminals, while the key buffer address field is used only for randomly accessed disk files.

\$RAB KRF

5.2.6 Key of Reference

The KRF parameter initializes the key of reference field, which specifies the key or index (primary, first alternate, and so on) to which the operation applies. The KRF field is applicable to indexed files only.

When your program issues a \$GET or \$FIND macro instruction in random access mode, the key of reference specifies the index to search for a match on the key value which is described by the key buffer address (KBF) and key size (KSZ) field. When your program issues a \$CONNECT

THE RECORD ACCESS BLOCK

or \$REWIND macro instruction, the key of reference identifies the index in the file of the next record in the stream. The next record is important in sequential retrieval of records; the Next Record is described in Section 10.2.2.

Format

KRF=key-number

key number

The numeric value representing a key in the records of a file. The value 0 indicates the primary key. The values 1 through 254 indicate alternate keys. The default value is 0 (primary key).

As an example, if the first alternate key is the index to search for a match (approximate, generic, or generic-approximate) on the key value described by the KBF and KSZ fields, the KRF parameter would be initialized as follows:

```
$RAB KRF=1,KBF=KEYBUF,KSZ=KEYSIZE
```

\$RAB KSZ

5.2.7 Key Size

The KSZ parameter initializes the key size field, which contains the size, in bytes, of the key pointed to by the key buffer address field.

Format

KSZ=size

size

The numeric value of the size of the record key. For relative record numbers, the default value of 0 causes a key size of 4 to be used. For string keys a value from 1 through the size of the key field and for the numeric key data types a value of 0 cause the defined size to be assumed; a nonzero value is checked against the defined size and an error (RMS\$_RSZ) is returned if they are not equal.

5.2.7.1 Relative Files - The size of the relative record number of a record in a relative file is a longword, positive, integer value; therefore, the key size is 4.

For example, for relative files the KSZ parameter must be coded as:

```
$RAB KSZ=4,KBF=RELKEY
```

5.2.7.2 Indexed Files. - The size of key values in bytes of an indexed file can be from 1 to 255 bytes.

When you access an indexed file in random mode, the contents of the KSZ and the contents of the ROP field determine the type of match to make on the key value specified in the key buffer address. For string

THE RECORD ACCESS BLOCK

key data type, the contents of the KSZ field can be less than the defined key size. For the other (numeric) key types, the contents of KSZ must be the defined length or 0, which defaults to the defined length. The following chart shows the relationships of the KSZ/ROP field contents and the type of match. Since KSZ for numeric key types must be the defined length, only exact and approximate matches are made on these types.

KGE or KGT Specified in ROP	Specified KSZ/Defined Key Size Relationship	Type of Match
NO	EQUAL	EXACT
NO	LESS THAN	GENERIC
YES	EQUAL	APPROXIMATE
YES	LESS THAN	GENERIC-APPROXIMATE

For example, the KSZ parameter for indexed files might be coded as follows to provide an approximate match on the first three characters:

```
$RAB KSZ=3,ROP=KGE
```

NOTE

The key size field uses the same location in the RAB as the prompt buffer size field. There is no conflict between these fields, however, because one field (PBF) is used only for terminal I/O, while the other field (KSZ) is used only for randomly accessed disk files.

\$RAB MBC

5.2.8 Multiblock Count

The MBC parameter initializes the multiblock count field, and applies only when the RAB accesses a sequential disk file.

VAX-11 RMS examines the multiblock count field during the execution of a \$CONNECT macro instruction. The value in this field is used as the number of blocks to be transferred as a single entity during an I/O operation for the record stream represented by this RAB. A buffer is allocated that can contain the specified number of blocks. In addition, more than one buffer (of this size) can be allocated for the record stream, as determined by the value of the multibuffer count field (see Section 5.2.9).

The use of the multiblock count field optimizes data throughput especially for sequential operations and in no way affects the structure of the file. It reduces the number of disk accesses you would normally require for your record operations and can thereby greatly increase execution speed. On the other hand, the extra buffering increases memory requirements.

THE RECORD ACCESS BLOCK

Format

MBC=blocks

blocks

The number of blocks, in the range of 1 through 127, to be allocated to each I/O buffer. If you omit this parameter, the multiblock count field is initialized to 0, which specifies that the process default for the multiblock count is to be used. If the process default is also 0, VAX-11 RMS uses the system default. If the system default is also 0, then the default size for each I/O buffer is one block. The DCL command SET RMS_DEFAULT is used to set process or system defaults.

For example, to allocate 16 blocks to each I/O buffer, the syntax is:

```
$RAB  MBC=16,CTX=RECOK
```

NOTE

The MBC parameter is not used with block I/O.

\$RAB MBF

5.2.9 Multibuffer Count

The MBF parameter sets the multibuffer count field to indicate the number of I/O buffers you want VAX-11 RMS to allocate when you issue a \$CONNECT macro instruction for this RAB.

VAX-11 RMS requires that at least one buffer be allocated for sequential and relative files and at least two buffers be allocated for indexed files, unless the file is to be processed with block I/O operations only. Multiple buffers can be used efficiently to overlap I/O time with program compute time, particularly in read-ahead or write-behind processing (see Section 5.2.15).

Format

MBF=buffers

buffers

A numeric value, in the range of -128 to +127, represents the number of buffers to be allocated. The absolute value of the field is used.

If the MBF parameter is omitted, the field is initialized to 0 at assembly time. A 0 value indicates the use of the process default for the particular file organization and device type.

If the process default is also 0, the system default for the particular file organization and device type applies.

If the system default is likewise 0, one buffer is allocated. However, if read-ahead or write-behind is specified at connect-time, a minimum of two buffers will be allocated. A minimum of two buffers will also be allocated for an indexed sequential file.

THE RECORD ACCESS BLOCK

For example, to allocate four buffers, the syntax is:

```
$RAB MBF=4,CTX=RECOK
```

NOTE

The MBF parameter is not used with block I/O. No buffers are allocated either if block I/O access is specified in the file access (FAC) field of the FAB on open or create, or if mixed block I/O and record I/O is specified in the file access field, but the block I/O record option is set in the record processing options (ROP) field for the connect service.

\$RAB PBF

5.2.10 Prompt Buffer Address

The PBF parameter initializes the prompt buffer address field. This field points to a character string to be used as a prompt for terminal input. If you select the PMT option of the ROP parameter (see Section 5.2.14) when you issue a \$GET macro instruction, this character string is output to the terminal before the read operation is performed.

To perform any carriage control on the terminal, you must insert the appropriate carriage control characters into this character string.

Format

```
PBF=prompt-address
```

prompt-address

The symbolic address of the buffer containing the prompt character string.

For example, if the buffer containing the prompt character string has a symbolic label of PROMPT, the PBF parameter is:

```
$RAB PBF=PROMPT,ROP=PMT,PSZ=2
```

NOTE

The prompt buffer address field uses the same location in the RAB as the key buffer address field. There is no conflict between these two fields, however, because the prompt buffer address field is used only for terminals, while the key buffer address field is used only for randomly accessed disk files.

\$RAB PSZ

5.2.11 Prompt Buffer Size

The PSZ parameter initializes the prompt buffer size field. This field contains the size, in bytes, of the character string for terminal I/O prompting.

Format

PSZ=prompt-size

prompt-size

The size, in bytes, of the prompt character string, in the range of 0 through 255.

If, for example, the character string is only two bytes long, the syntax is:

\$RAB PBF=PROMPT,PSZ=2,ROP=PMT

NOTE

The prompt buffer size field uses the same location in the RAB as the key size field. There is no conflict between these fields, however, because the prompt buffer size (PSZ) field is used only for terminal I/O, while the key size (KSZ) is used only for randomly accessed disk files.

\$RAB RAC

5.2.12 Record Access Mode

The RAC parameter initializes the record access mode field to indicate the method of retrieving or storing records in the file.

Format

RAC= $\left\{ \begin{array}{l} \text{SEQ} \\ \text{KEY} \\ \text{RFA} \end{array} \right\}$

SEQ

Indicates sequential record access mode (the default); can be specified with any type of file organization.

KEY

Indicates random access by key; used with relative files (and with sequential files on disk with fixed-length records) to indicate access by relative record number; used with indexed files to indicate access by key value.

THE RECORD ACCESS BLOCK

RFA

Indicates random access by record's file address; used for disk files only.

For example, to set the record access mode field to indicate the sequential record access mode, the syntax is:

```
$RAB RAC=SEQ,CTX=RECOK
```

The offset for this field is:

```
RAB$B_RAC
```

Each record access mode has its own symbolic value.

- SEQ - RAB\$C_SEQ
- KEY - RAB\$C_KEY
- RFA - RAB\$C_RFA

NOTES

1. You can specify the record access mode on a per-operation basis.
2. For block I/O, you do not use the record access mode field.

\$RAB RBF

5.2.13 Record Address

The RBF parameter initializes the record address field. When you issue a \$PUT or \$WRITE macro instruction, this field must specify the address of the record to be written to the file.

When you issue a \$GET or \$READ macro instruction, VAX-11 RMS sets this field to the address of the record just read from the file; you need not initialize this field.

Format

```
RBF=buffer-address
```

buffer-address

The symbolic address of the buffer in your program that contains the record to be written.

For example, to initialize the record address field with the address of a buffer having the label of RECBUF, the syntax is:

```
$RAB RBF=RECBUF,CTX=RECOK
```

\$RAB RHB**5.2.14 Record Header Buffer**

The RHB parameter initializes the fixed-length record header field. This buffer is used only when processing records of variable with fixed-length control. For a \$GET macro instruction, VAX-11 RMS strips the fixed control area portion of the record and places it in the buffer whose address is specified in this field. For the \$PUT or \$UPDATE macro instructions, VAX-11 RMS writes the contents of the specified buffer to the file as the fixed control area portion of the record.

The size of this fixed control area is defined in the FAB, through the FSZ parameter. You must ensure that the size of the buffer described in the record header buffer field is equal to the value specified by the FSZ parameter.

Format

RHB=header-address

header-address

The symbolic address of the record header buffer. If omitted, an address of 0 is assumed, which indicates the absence of a buffer; the fixed control area is discarded for a \$GET macro instruction, zeroed for a \$PUT macro instruction, and left unchanged for a \$UPDATE macro instruction.

For example, if the buffer is defined with a label of FCABUF, the syntax is:

\$RAB RHB=FCABUF,CTX=RECOK

\$RAB ROP**5.2.15 Record-Processing Options**

The ROP parameter sets indicators in the record-processing options field that let you request optional functions during execution of a record operation. VAX-11 RMS operations never modify the contents of this field.

Format

ROP= <ASY,BIO,CCO,CVT,EOF,LOC,KGE,KGT,LOA,LIM,NLK,NXR,PMT,PTA,
RAH,RLK,RNE,RNF,TMO,TPT,UIF,ULK,WBH>

Options that are input to \$CONNECT:

ASY

Asynchronous: See detailed explanation below. Please note that for indexed files, I/O may take place during the \$CONNECT.

THE RECORD ACCESS BLOCK

BIO

Block I/O: This option is meaningful only if the BRO (FOP field in the FAB) was set on \$OPEN or \$CREATE. Setting BIO on \$CONNECT declares that only block I/O operations will be permitted. If BIO is clear on \$CONNECT, only record operations will be allowed for relative and indexed files, or mixed operations will be allowed on sequential files. See Chapter 12 for further discussion.

EOF

End-of-file; indicates that VAX-11 RMS is to position to the end of the file when a \$CONNECT macro instruction executes. This applies only to sequential disk files.

RAH and WBH

Read ahead and Write behind: If either the RAH or WBH is set, and the multibuffer count (see MBF in the RAB) is 0, two buffers will be allocated to allow multibuffering. If two or more buffers are specified, multibuffering will be allowed regardless of the setting at \$CONNECT. Conversely, if a buffer count of 1 is specified, multibuffering is disabled regardless of the setting at \$CONNECT.

Options applicable to indexed sequential files only:

These options are selectable on a per-operation basis, i.e., they may be enabled or disabled on any operation.

KGE

Key is greater than or equal to; requests VAX-11 RMS to access the first record in an indexed file, which contains a value for the specified key of reference (KRF) (see Section 5.2.6) that is greater than or equal to the value described by the key buffer address (KBF) and key size (KSZ) fields (see Section 5.2.5 and 5.2.7.2, respectively). If neither KGE nor KGT is specified, a key equal match is made.

KGT

Key is greater than; requests VAX-11 RMS to access the first record in an indexed file, which contains a value for the specified key of reference (KRF) (see Section 5.2.) that is greater than the value described by the key buffer address (KBF) and key size (KSZ) fields (see Sections 5.2.5 and 5.2.7.2, respectively). If neither KGE nor KGT is specified, a key equal match is made.

LOA

Load; specifies that VAX-11 RMS is to load buckets according to the fill size established at file creation time. The bucket fill size is established at file creation time by the data bucket fill size (DFL) and index bucket fill size (IFL) fields of the key extended attribute blocks (XABs). The XABs are described in Chapter 6. If LOA is not specified, VAX-11 RMS ignores the established bucket fill size (that is, buckets will be completely filled).

THE RECORD ACCESS BLOCK

LIM

Limit; the key value described by the key buffer address (KBF) and key size (KSZ) fields (see Sections 5.2.5 and 5.2.7.2, respectively) is to be compared to the value in the record accessed in sequential mode. If the record's key value is greater than the limit key value, and RMS\$_OK_LIM status code is returned.

Options affecting record operation performance:

These options are selectable on a per-operation basis, i.e., they may be enabled or disabled on any operation.

ASY

Asynchronous; indicates that this I/O operation is to be performed asynchronously. When you specify ASY, VAX-11 RMS will return control to your program as soon as an I/O operation is initiated, even though that operation may not yet be completed. This is normally used in conjunction with the \$WAIT macro to synchronize with operation completion. See Chapter 10 for further discussion.

FDL

Fast delete: This applies only to \$DELETE operation on indexed sequential files. When specified, the pointers from alternative indices which allow duplicates are not removed. This saves an index search on those indices when deleting records.

LOC

Locate mode; indicates that record operations involving the \$GET macro instruction will use locate mode (see Section 10.1.2).

RAH

Read-ahead; used with multiple buffers (see Section 5.2.8) to indicate read-ahead operations. When a buffer is filled, the next record will be read into the next buffer. This permits an overlapping of input and computing. Read-ahead is ignored for unit record device I/O. This option is implemented only for the sequential file organization.

WBH

Write-behind; used with multiple buffers (see Section 5.2.8). When a buffer is filled, the next record written will be placed in the next buffer while the previous buffer is output. This allows for an overlapping of computing and output. Write-behind is ignored for unit record devices. This option is implemented only for the sequential file organization.

Options controlling record locking:

These options apply only to relative, indexed, and sequential files with 512-byte fixed-length records. These options are selectable on a per-operation basis, i.e., they may be enabled or disabled on any operation.

THE RECORD ACCESS BLOCK

NLK

No lock; specifies that the record accessed through a \$GET or \$FIND macro instruction is not to be locked. The NLK option takes precedence over the ULK option (below).

NXR

Nonexistent record processing; specifies that if the record randomly accessed through a \$GET or \$FIND macro instruction does not exist (was never inserted into the file or was deleted), the service is to be performed anyway, locking the record cell if required. For the \$GET macro instruction, the previous contents of a deleted record are returned. The processing of a deleted record returns a completion status code of RMS\$OK_DEL, and the processing of a record that never existed returns RMS\$OK_RNF. This option does not apply to indexed sequential files.

RLK

Read of locked record allowed; specifies that a user who locks a record is allowing the locked record to be read by other accessors.

ULK

Manual unlocking; specifies that VAX-11 RMS cannot automatically unlock records. Instead, once locked (through a \$GET, \$FIND or \$PUT macro instruction), a record must be specifically unlocked by a \$FREE or \$RELEASE macro instruction. The NLK option (above) takes precedence over the ULK option.

Options relevant to Put operations only:

These options are selectable on a per-operation basis, i.e., they may be enabled or disabled on any operation.

TPT

Truncate put; specifies that a put service with a record access mode of sequential can occur at any point in the file, truncating the file at that point. On a write service, this causes the end of file mark to immediately follow the last byte written. This applies only to sequential files.

UTF

Update if; indicates that if a \$PUT macro instruction is issued for a record that already exists in the file, the operation is converted to an update. This option is necessary to overwrite (as opposed to update) an existing record in relative and indexed sequential files. Indexed files using this option, must not allow duplicates on the primary key.

Miscellaneous options:

These options are selectable on a per-operation basis, i.e., they may be enabled or disabled on any operation.

BIO

Block I/O; this option may also be used to mix block and record operations to sequential files. See Chapter 12.

THE RECORD ACCESS BLOCK

TMO

Timeout; in addition to its use for terminals, the TMO option serves a special purpose for mailbox devices. If specified along with a time-out value of zero (TMO field in the RAB), \$GET and \$PUT operations to mailboxes will use the IO\$M NOW modifier. This will cause the operation to complete immediately, instead of synchronizing with another cooperating writer or reader of the mailbox. See the VAX-11 I/O User's Guide for a further discussion of mailboxes.

Options specific to terminal devices.

These options map directly into equivalent modifiers in the QIO function code. For a further discussion of their effects, see the VAX/VMS I/O User's Guide. These options are selectable on a per-operation basis, i.e., they may be enabled or disabled on any operation.

CCO

Cancel control 0; guarantees that terminal output will not be discarded if the operator has entered CTRL/O.

CVT

Convert; changes characters to uppercase on a read from a terminal.

PMT

Prompt; indicates that the contents of the prompt buffer are to be used as a prompt on a read from a terminal (see Section 5.2.10).

PTA

Purge type-ahead; eliminates any information that may be in the type-ahead buffer on a read from a terminal.

RNE

Read no echo; indicates that input data is not echoed (displayed) on the terminal as it is entered on the keyboard.

RNF

Read no filter; indicates that CTRL/U, CTRL/R, and DELETE are not to be considered control commands on terminal input, but are to be passed to the user program.

TMO

Time-out; indicates that the content of the time-out period field of the RAB is to be used to determine the number of seconds that a VAX-11 RMS operation has to complete its operation. If the time-out period expires, VAX-11 RMS returns an error status (see Section 5.2.17).

THE RECORD ACCESS BLOCK

You can use one or more options with the ROP parameter. For example, to indicate that a terminal read should convert from lower- to uppercase, and use locate mode, the prompt buffer, and the specified time-out period, the ROP parameter would be:

```
$RAB ROP=<CVT,LOC,PMT,TMO>,PBF=PROMPT,PSZ=PROMPT_SIZE,TMO=30
```

Each option has its own symbolic bit offset and mask value.

\$RAB RSZ

5.2.16 Record Size

The RSZ parameter sets the record size field. This field controls the size of a record or the number of bytes that, respectively, a \$PUT or \$WRITE (block I/O) macro instruction can write.

On input from a file, VAX-11 RMS sets this field to indicate the length, in bytes, of the record that a \$GET macro instruction transfers or that a \$READ macro instruction reads.

Format

```
RSZ=record-size
```

record-size

The size, in bytes, of the record. For operations with a \$WRITE macro instruction, the range is 1 through 65535. \$PUT operations may specify a size from 0 to the maximum shown in the following table:

FILE ORGANIZATION	RECORD FORMAT	MAXIMUM ALLOWED
Sequential	Fixed-length	32,767
Sequential (Disk)	Variable-length	32,767-FSZ ¹
Sequential (ANSI Tape)	Variable-length	9,995-FSZ
Relative	Fixed-length	16,383
Relative	Variable-length	16,381-FSZ
Indexed Sequential	Fixed-length	16,362
Indexed Sequential	Variable-length	16,360

For example, to indicate a record size of 150 bytes, the syntax is:

```
$RAB RBF=RECBUF,RSZ=150
```

1. The FSZ represents the size of the fixed control area of a record. The FSZ is equal to 0 for variable-length records. The FSZ is equal to the size in bytes, for the fixed control area of the VFC (variable with fixed-length control) records.

THE RECORD ACCESS BLOCK

NOTES

1. After a get operation, VAX-11 RMS places the size of the record retrieved into the record size field. On a read operation, VAX-11 RMS sets the record size field to the number of bytes actually transferred.
2. For variable with fixed-length control records, VAX-11 RMS does not include the size of the fixed control area in the record size field.

\$RAB TMO

5.2.17 Time-Out Period

The TMO parameter initializes the time-out period field, which indicates the maximum number of seconds that VAX-11 RMS can use to complete an operation. If the time-out period expires before the operation completes, VAX-11 RMS returns an error status code.

To use this field, you must also specify the TMO option when you set the record-processing option field (ROP parameter).

Format

TMO=seconds

seconds

The maximum number of seconds, in the range of 0 through 255, that a \$GET from the terminal operation can use. If you specify 0, the current contents of the type ahead buffer is returned.

For example, to indicate that a \$GET for a terminal must complete in 20 seconds or less, the syntax is:

```
$RAB TMO=20,ROP=TMO
```

Note that the TMO option must also be specified on the ROP parameter.

NOTE

A TMO of 0 for either a \$GET or \$PUT to a mailbox will cause the operation to complete immediately, rather than waiting for another process. For example, a \$PUT with a TMO field of 0, to a mailbox device will not wait for another process to read the record.

\$RAB UBF

5.2.18 User Record Area Address

The UBF parameter initializes the user record area address field, which indicates the location of a record or block buffer.

When you issue a \$GET macro instruction, this field must contain the buffer address regardless of the record transfer mode (locate or move). This option also applies when you issue a \$READ macro instruction for block I/O. However, operations with a \$PUT macro instruction never need a user buffer.

Format

UBF=buffer-address

buffer-address

The symbolic address of a work area (buffer) within your program. (The size of this buffer must be defined in the user record area size field; the USZ parameter.)

For example, if the buffer area has a label of USRBUF, the syntax is:

```
$RAB UBF=USRBUF,USZ=2048
```

\$RAB USZ

5.2.19 User Record Area Size

The USZ parameter initializes the user record area size field, which indicates the length, in bytes, of the user record or block buffer. This buffer area should be large enough to contain the largest record in the file. If the buffer is not large enough on an operation with a \$GET macro instruction, VAX-11 RMS will move as much of the record as possible into the buffer, and return a warning status code.

The value in this field specifies the transfer length, in bytes, for block I/O operations with a \$READ macro instruction.

Format

USZ=buffer-size

buffer-size

A numeric value representing the size, in bytes, of the buffer. This value must be in the range of 1 through 65535.

For example, for a user buffer area with a label of USRBUF and a size of 2048 bytes, the syntax is:

```
$RAB UBF=USRBUF,USZ=2048
```

THE RECORD ACCESS BLOCK

5.3 NONINITIALIZABLE RAB FIELDS

The following list describes the RAB fields that you cannot initialize at assembly time. Either they are static, or VAX-11 RMS sets them for you.

Operation Completion Status Fields:

STS

Completion status code field; VAX-11 RMS sets this field with the success or failure status codes for a record operation before returning control to your program. In the case of an asynchronous operation that has been initiated but not yet completed, this field is 0. When the operation is complete, the field will be updated with the completion status. See Section 8.4 for additional details about RMS status codes. Potential error codes for specific operations are listed with their description in Chapters 11 and 12. Appendix A lists the symbolic completion status codes that your program can use to test the contents of this field.

STV

Status value field; communicates additional completion information to your program, based on the type of operation and the contents of the completion status code field. For additional information on the STS and STV fields, see Section 8.4. See Appendix A for the instances when VAX-11 RMS uses the status value field.

Internal Stream Identifier Field:

ISI

Internal stream identifier field; associates the RAB with a corresponding FAB. VAX-11 RMS sets this field after the execution of a \$CONNECT macro instruction. A \$DISCONNECT macro instruction clears this field. This field should not be altered.

Static Fields:

BID

Block identifier field; identifies the block as a RAB. The \$RAB macro instruction sets this field to the symbolic value RAB\$C_BID; this field must not be altered.

BLN

Block length field; defines the length, in bytes, of the RAB to VAX-11 RMS. The \$RAB macro instruction sets this field to the symbolic value RAB\$C_BLN; this field must not be altered.

5.3.1 The Record's File Address

After the successful execution of a \$GET, \$PUT, or \$FIND macro instruction, VAX-11 RMS sets the record's file address (RFA) field to the address of the record acted on by the operation. This address is meaningful only for disk files; it provides an unambiguous means of randomly locating this same record at some later time.

THE RECORD ACCESS BLOCK

You can store the contents of the record's file address field for future use. When you want to retrieve the record again, merely restore the saved contents of the field, set the record access mode to random by RFA, and issue a \$GET or \$FIND macro instruction.

NOTES

1. This field is six bytes long. There are two ways to refer to this field:

- (1) RAB\$W RFA is the offset for this field.
RAB\$\$ RFA is the size of this field

The field may be copied:

```
MOVAL    RABBLK,R0
MOV3     #RAB$$ RFA,-
         RAB$W RFA (R0),-
         SAVE RFA
```

- (2) There are two additional names for portions of this field:

RAB\$L RFA0 is the offset of the first longword
RAB\$W RFA4 is the offset of the last word

The field may be copied:

```
MOVAL    RABBLK,R0
MOVL     RAB$L RFA0 (R0),-
         SAVE RFA
MOVW     RAB$W RFA4 (R0),
         SAVE RFA+4
```

2. RFA values remain valid for a record in a sequential file as long as the record is within the space defined by the logical file, that is, until the file is truncated to a point before the record.
3. RFA values remain valid for a record in a relative or indexed file for the life of the file, that is, until the file is deleted.

CHAPTER 6

THE EXTENDED ATTRIBUTE BLOCKS

This chapter describes the various Extended Attribute Blocks (XABs), their fields, and the macro instructions and parameters you use to initialize the fields at assembly time.

6.1 THE PURPOSE OF EXTENDED ATTRIBUTE BLOCKS

The XABs are optional additional control blocks, which you can use to communicate to VAX-11 RMS any file attributes beyond those expressed in the FAB. You use these control blocks only when you want to specify exactly, or retrieve information on, the attributes handled by a particular XAB.

You can use XABs to set file attributes by specifying them as inputs to the \$CREATE, \$CLOSE, or \$EXTEND macro instructions. Retrieve the attributes by specifying the XAB as input to the \$OPEN or \$DISPLAY macro instructions. If the Create-if (CIF) bit is set in the file processing options field of the FAB on a create service, VAX-11 RMS uses the XAB fields as input or output depending on whether the file is opened or created, respectively.

When you need more than one XAB, you can chain them together. Each XAB has a next XAB address field, which can be set at assembly time through the NXT parameter, or at run time. You can set this field at run time by storing the appropriate address into the next XAB address field. The extended attribute block pointer field of the FAB (see Section 4.2.24) points to the first block in the chain. Section 6.2 below describes chaining in detail.

Currently, VAX-11 RMS supports seven types of XABs, each with its own macro instructions for allocation and initialization. These blocks and their macro instructions are as follows:

- Allocation control -- \$XABALL
- Date and time -- \$XABDAT
- File header characteristics -- \$XABFHC
- File protection -- \$XABPRO
- Key definition -- \$XABKEY
- Revision date and time -- \$XABRDT
- Summary -- \$XABSUM

THE EXTENDED ATTRIBUTE BLOCKS

The last three characters of each macro instruction (ALL, DAT, FHC, KEY, PRO, RDT, SUM) define the specific type of the XAB to VAX-11 RMS, and cause the value for this specific type to be stored in the type code field of each block. The symbolic offset for this field is:

XAB\$B_COD

The symbolic values stored in the type code field are:

Allocation control	-- XAB\$C_ALL
Date and time	-- XAB\$C_DAT
File header characteristics	-- XAB\$C_FHC
File protection	-- XAB\$C_PRO
Key definition	-- XAB\$C_KEY
Revision date and time	-- XAB\$C_RDT
Summary	-- XAB\$C_SUM

In addition, a length value is stored in the block length field of each block. The symbolic offset for this field is:

XAB\$B_BLN

The symbolic values stored in the block length field are:

Allocation control	-- XAB\$C_ALLEN
Date and time	-- XAB\$C_DATLEN
File header characteristics	-- XAB\$C_FHCLEN
File protection	-- XAB\$C_PROLEN
Key definition	-- XAB\$C_KEYLEN
Revision date and time	-- XAB\$C_RDTLEN
Summary	-- XAB\$C_SUMLEN

Because each block has its own initialization macro instruction, each is discussed separately in Sections 6.3 through 6.9 below. Each XAB macro instruction may be prefixed by an optional label. This label lets you assign a name to an XAB, thereby allowing symbolic access to the XAB.

For example, suppose a \$XABDAT macro instruction is used and has the following label:

```
DATE_XAB:    $XABDAT
```

Then, your \$FAB macro instruction would have an XAB parameter as follows:

```
$FAB      XAB=DATE_XAB
```

Note that the label must be separated from the XAB macro instruction by a colon (:).

Table 6-1 indicates which XAB types are processed by which service.

THE EXTENDED ATTRIBUTE BLOCKS

Table 6-1
XAB Types Processed by Service

Service Type	Close	Create	Display	Extend	Open
Allocation Control		Input Output	Output	Input Output	Output
Key Definition		Input Output	Output		Output
Summary		Output	Output		Output
Date and Time		Input Output ¹	Output		Output
File Header Characteristics		Input Output ¹	Output		Output
File Protection	Input ²	Input Output ¹	Output		Output
Revision Date and Time	Input ²	Input Output ¹	Output		Output

¹ Fields of the XAB are output only if the create if (CIF) bit is set and the file is opened, not created.

² Processed only if file is write-accessed.

At assembly time, you can initialize the fields of the particular XAB through keyword parameters. At run time, you can use the keyword parameters with the appropriate \$XABxxx_STORE macro (see Chapter 14) or the defined symbolic offsets.

6.2 CHAINING EXTENDED ATTRIBUTE BLOCKS

Every XAB has a next XAB address field, regardless of the type of information that the XAB contains, such as date/time or file protection information. When you need one or more XABs for a particular operation, place the symbolic address of the first XAB of the chain into the extended attribute block pointer field of the FAB. Then, place the address of the second XAB in the chain, if there is one, in the next XAB address field of the first XAB (NXT parameter). Continue this process until you have chained all the XABs you need. You must set the next XAB address field of the last XAB to 0 to indicate the end of the chain. You can either set this field explicitly or allow the system to default to the 0 value.

Within the XAB chain, the different types of XABs need not be in any specific order. For example, at assembly time you could allocate a date and time XAB, a file protection XAB, and an allocation control XAB. You can chain these different types of XABs in any order by appropriately setting the contents of the next XAB address field in each block.

THE EXTENDED ATTRIBUTE BLOCKS

For indexed files, however, VAX-11 RMS permits multiple instances of the same type of XAB in an allocation control or key definition XAB chain. For \$CREATE macro instructions, the multiple instances must appear in a specific order and allocation control XABs must be linked together in ascending order based on the contents of the area identification number (AID) field (see Section 6.5.2); key definition XABs must be linked together in ascending order based on the contents of the key of reference (REF) field (see Section 6.6.12). Also, for \$CREATE macro instruction, there cannot be any intervening XABs of another type in the subchain of XABs of one type.

Further, the operation for which the allocation control or key definition XABs is present determines whether the ascending order must be dense. For create operations, allocation control and key definition XABs, if present, must appear in densely ascending order by area identification (AID) number or key of reference (REF) value, respectively. For extend operations, allocation control and key definition XABs, if present, must be in ascending order but need not be dense. For open and display operations, RMS-32 verifies that the number of XABs specified does not exceed the number specified for the file. If the number of XABs specified does exceed the number defined for the file, a RMS\$_AID error is returned for allocation XABs and a RMS\$_REF error is returned for key definition XABs.

The NXT parameter appears in the format of each of the XAB macro instructions. This parameter is explained below, rather than repeated throughout Sections 6.3 through 6.10.

Format

NXT=address

address

The symbolic address of the next XAB in the chain. A value of 0 (the default) indicates the last (or only) XAB in the chain.

\$XABDAT

6.3 DATE AND TIME XAB

The \$XABDAT macro instruction allocates and initializes an XAB for date and time. This block allows for extended control of the date and time of the file's creation, revision (update), and expiration. Table 6-2 summarizes the fields comprising the date and time XAB.

THE EXTENDED ATTRIBUTE BLOCKS

Table 6-2
Date and Time Extended Attribute Block Fields

Field Name	Field Size	Description	Offset
BLN ²	byte	Block length	XAB\$B_BLN
CDT ¹	quadword	Creation date and time	XAB\$Q_CDT
COD ²	byte	Type code	XAB\$B_COD
EDT	quadword	Expiration date and time	XAB\$Q_EDT
NXT	longword	Next XAB address	XAB\$L_NXT
RDT ¹	quadword	Revision date and time	XAB\$Q_RDT
RVN ¹	word	Revision number	XAB\$W_RVN

¹ Indicates no assembly time initialization.

² Indicates that this field is set automatically by the type of macro instruction.

Format

OPERATION	PARAMETERS
label: \$XABDAT	EDT=date-time NXT=address

\$XABDAT EDT

6.3.1 Expiration Date and Time

The EDT parameter sets the expiration date and time field. This field indicates the date and time after which a magnetic tape file can be deleted. It is not currently used for disk files, and its future use is reserved to DIGITAL.

Format

EDT=date-time

date-time

A 64-bit binary value in either absolute (positive) or delta (negative) format. (See the VAX/VMS System Services Reference Manual.)

THE EXTENDED ATTRIBUTE BLOCKS

6.3.2 Creation/Revision Date and Time, and Revision Number

VAX-11 RMS sets certain values for date and time, and returns them in date and time XAB fields for your inspection. You can override these system-supplied values through the use of a date and time XAB as input to a \$CREATE macro instruction. However, the \$XABDAT macro instruction does not contain parameters for the assembly-time initialization of these fields. As outlined in Table 6-2, these fields are:

- Creation date and time (CDT) -- this is a 64-bit binary value expressing the date and time at which the file was created.
- Revision date and time (RDT) -- this is a 64-bit binary value expressing the date and time at which the file was last updated.
- Revision Number (RVN) -- this field provides the number of times this file was opened for write operations.

The following table describes how the fields of the XABDAT block are used by the RMS file-processing macro instructions.

OPERATION	CDT	EDT	RDT
CLOSE	not used	not used	not used
CREATE	INPUT	INPUT	INPUT 1
DISPLAY	OUTPUT	OUTPUT	OUTPUT
ERASE	not used	not used	not used
EXTEND	not used	not used	not used
OPEN	OUTPUT	OUTPUT	OUTPUT

NOTE

If the user specifies the CDT or RDT field in the XABDAT as zero (either explicitly or by default) and, if the specified field is used by the file-processing macro, it will be replaced with the current date and time.

NOTE

If the CREATE macro is invoked with a FAB that has the CIF (create if) bit set and the file to be created already exists, the CREATE is processed like an OPEN and the fields listed above are outputs.

1. The RDT field in CREATE is superseded by the current date and time on CLOSE. In order to specify this field, the user should employ the XABORT block.

\$XABPRO

6.4 FILE PROTECTION XAB

The \$XABPRO macro instruction allocates and initializes an XAB that you can use to explicitly specify file ownership and file protection. Table 6-3 summarizes the fields comprising the file protection XAB.

Table 6-3
File Protection Extended Attribute Block Fields

Field Name	Field Size	Description	Offset
BLN ¹	byte	Block length	XAB\$B_BLN
COD ¹	byte	Type code	XAB\$B_COD
GRP	word	Group number of file owner	XAB\$W_GRP
MBM	word	Member number of file owner	XAB\$W_MBM
NXT	longword	Next XAB address	XAB\$L_NXT
PRO	word	File protection; contains four separate fields denoting the protection for system, owner, group, and world	XAB\$W_PRO

¹Indicates that this field is set automatically by the type of macro instruction.

Format

OPERATION	PARAMETERS
label: \$XABPRO	PRO=<system,owner,group,world> UIC=<group,member> NXT=address

\$XABPRO PRO

6.4.1 File Protection

The PRO parameter initializes the four subfields of the file protection field and it specifies the file access privileges of the four classes of users. The subfields for the four classes are:

1. System -- specifies access rights for users executing under a system UIC, that is, having a group number less than 8.
2. Owner -- specifies access rights for the owner of the file. A user is considered the owner of the file only if both the group and member number fields (see Section 6.4.2) of the accessing process match the group and member number fields of the file owner's UIC stored with the file.
3. Group -- specifies the access rights for users whose group number matches the group number field of the file owner.
4. World -- specifies the access rights for any user. It is normally allowed for users not within the system, owner, or group classifications (items 1, 2, and 3, above).

A user is granted the maximum number of types of access rights for each of the classes to which he belongs.

The entire file protection field is one word, and each classification subfield occupies four bits of this word. The field is organized as shown in Figure 6-1.

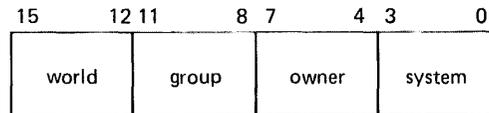


Figure 6-1 File Protection Field

Format

PRO=<SYSTEM,OWNER,GROUP,WORLD>

<SYSTEM,OWNER,GROUP,WORLD>

The access code for the four classifications of users. An access code consists of four bits, each of which represents the type of access granted to a user in the class. These access rights and the characters that signify them are:

- R - read access
- W - write access
- E - execute access
- D - delete access

You can specify any number of access characters, in any order, for each classification. For example, you could specify RWD, RWED, DREW, or any combination, up to four characters per classification. The

THE EXTENDED ATTRIBUTE BLOCKS

access rights for one classification are not separated by a comma. However, the classifications must be separated from each other by a comma or other valid separator to delimit the end of one classification and the start of the next. For example, the access rights for one classification may have a syntax of:

```
RWD
```

However, the syntax for three separate classifications might be:

```
RWD,DWRE,R
```

Note that when you use less than all four access rights characters, you need not supply a delimiter or code to indicate the omission.

The four different classifications of users, however, must be coded in the following order:

```
<SYSTEM,OWNER,GROUP,WORLD>
```

The angle brackets are required syntax, and each classification must be separated from the others by a comma. In addition, when you omit a classification, the comma must be retained to indicate the omission, unless no other classification follows. For example, to specify all access rights for system, owner, and world, you would write:

```
$XABPRO PRO=<RWED,RWED,,RWED>
```

However, to specify all access rights to only system and owner, you would write:

```
$XABPRO PRO=<RWED,RWED>
```

The absence of a code specifies that the access associated with the code is denied to the user.

Each 4-bit subfield also has its own symbolic offset, as follows:

- System -- XAB\$V_SYS
- Owner -- XAB\$V_OWN
- Group -- XAB\$V_GRP
- World -- XAB\$V_WLD

Additionally, each separate access specification has the following mask values:

- No read access -- XAB\$M_NOREAD
- No write access -- XAB\$M_NOWRITE
- No execute access -- XAB\$M_NOEXE
- No delete access -- XAB\$M_NODEL

THE EXTENDED ATTRIBUTE BLOCKS

User Consideration

The bit values in the protection world are set to 1 to deny access. Thus, specifying a particular access right code clears the bit to 0.

NOTE

If you do not provide a file protection XAB for a \$CREATE macro instruction, or if the PRO parameter is not specified or is specified as no access to all classes (all 1 bits), the default file protection for the process will be used for the newly created file.

\$XABPRO UIC

6.4.2 Group and Member Number

The UIC parameter initializes both the group and member number fields, thus supplying both portions of the user identification code (UIC) of the file's owner.

Format

UIC=<GROUP,MEMBER>

<GROUP,MEMBER>

The group number and member number, respectively, of the owner of the file. Both numbers are octal numbers in the range of 0 through 177777. The group number and member number must be enclosed within angle brackets, placed in the order shown in the format, and separated by a comma.

For example, if your group number is 126 and your member number is 1, the syntax is:

```
$XABPRO UIC=<126,1>
```

The symbolic offsets for these fields are:

- Group number -- XAB\$W_GRP
- Member number -- XAB\$W_MBM

The total user identification field, including both the group and member number fields, has a symbolic offset of:

```
XAB$L_UIC
```

NOTE

If no file protection XAB is provided, or the user identification field is null for a \$CREATE macro instruction, the UIC of the process will be used as the owner's UIC for the newly created file.

\$XABALL

6.5 ALLOCATION CONTROL XAB

The \$XABALL macro instruction allocates and initializes an XAB that allows extended control of file disk space allocation, both for initial allocation and later extension. When you use an allocation control XAB as input to a create or extend service, certain fields override corresponding fields of the FAB. Overriding occurs in the allocation quantity (ALQ), bucket size (BKZ), which is the BKS in the FAB), and default extension quantity (DEQ) fields, and in the CBT and CTG bits of the allocation options (AOP, which are the CBT and CTG bits of the FOP field in the FAB) field. On an open or display service, VAX-11 RMS fills in these fields with the values that pertain to the file. Table 6-4 summarizes the fields comprising the allocation control XAB.

Table 6-4
Allocation Control Extended Attribute Block Fields

Field Name	Field Size	Description	Offset
AID	byte	Area identification number	XAB\$B_AID
ALN	byte	Alignment boundary type	XAB\$B_ALN
ALQ	longword	Allocation quantity	XAB\$L_ALQ
AOP	byte	Allocation options	XAB\$B_AOP
BKZ	byte	Bucket size	XAB\$B_BKZ
BLN ¹	byte	Block length	XAB\$B_BLN
COD ¹	byte	Type code	XAB\$B_COD
DEQ	word	Default extension quantity	XAB\$W_DEQ
LOC	longword	Location	XAB\$L_LOC
NXT	longword	Next XAB address	XAB\$L_NXT
VOL	word	Relative volume number	XAB\$W_VOL

¹Indicates that this field is set automatically by the type of macro instruction.

THE EXTENDED ATTRIBUTE BLOCKS

Format

OPERATION	PARAMETERS
label: \$XABALL	AID=area-number ALN= $\left\{ \begin{array}{l} \text{CYL} \\ \text{LBN} \\ \text{VBN} \\ \text{RFI} \end{array} \right\}$ ALQ=allocation-qty AOP=<CBT, CGT, HRD, ONC> BKZ=bucket-size DEQ=extension-qty LOC=number NXT=address RFI=<F ₁ , F ₂ , F ₃ > VOL=volume-number

\$XABALL AID

6.5.1 Area Identification Number

The AID parameter initializes the area identification number field, which identifies the area of the file described by the current XAB. You are always responsible for the contents of this field; it is never set by VAX-11 RMS. VAX-11 RMS uses the contents of this field for the following:

- Checks the sequencing of allocation control XABs in an XAB chain. The allocation XABs in an XAB chain must appear in ascending order, based on the contents of the AID field in \$CREATE and \$EXTEND macro instructions; the order is irrelevant in the \$OPEN and \$DISPLAY macro instructions.
- Identifies the target area for a specific operation (for example, create, extend, and so on).

Format

AID=area-number

area-number

A numeric value indicating which area, in a range of 0 through 254, of the file is described by the current XAB. If the file is a sequential or relative file, only a single allocation XAB can be used for any operation and its AID field must contain 0. The default for this field is 0.

For example, to establish an allocation XAB for area 3 of an indexed file, you would write:

```
$XABALL AID=3
```

\$XABALL ALN**6.5.2 Alignment Boundary Type**

The ALN parameter initializes the boundary type field, which specifies the type of alignment for the area to be allocated. This gives you control over the placement of your file. If you need this placement control on either a create or extend operation, you use the alignment boundary type field to specify whether the location field (LOC parameter) contains a starting cylinder number, logical block number, or virtual block number.

Format

$$ALN = \left\{ \begin{array}{l} CYL \\ LBN \\ RFI \\ VBN \end{array} \right\}$$

CYL

Indicates that the alignment starts at the cylinder number specified in the location field.

LBN

Indicates that the alignment starts at the logical block number specified in the location field.

RFI

Indicates that the alignment starts as near as possible to the file specified by the related file identification field, at the virtual block number specified in the location field.

VBN

Indicates that the alignment starts as near as possible to the virtual block number specified in the location field.

For example, if you want the file you are going to create or extend to be aligned at the tenth cylinder on the volume, you would write:

```
$XABALL  ALN=CYL,LOC=10
```

Each alignment type has its own symbolic value.

- CYL - XAB\$C_CYL
- LBN - XAB\$C_LBN
- RFI - XAB\$C_RFI
- VBN - XAB\$C_VBN

NOTE

If you do not set a value in this field, VAX-11 RMS assumes that you do not want to exercise control over the placement of your file.

\$XABALL ALQ

6.5.3 Allocation Quantity

The ALQ parameter sets the allocation quantity field. This field indicates the number of blocks to be allocated initially, when using the \$CREATE macro instruction. It is also used to specify the number of blocks to add to the file, when using the \$EXTEND macro instruction. This parameter may be specified for each area in an indexed sequential file having multiple areas.

In either case (Create or Extend operation), the value in this field overrides the contents of the allocation quantity field of the FAB (see Section 4.2.2).

The Open, Create, and Display services fill in this field with the actual allocation size of the file or area for indexed files. The extend service fills in the field with the actual size of the extended space.

Format

ALQ=allocation-quantity

allocation-quantity

A numeric value in the range of 0 through 4,294,967,295. A value of 0 (the default) indicates that no allocation is to be performed.

For example, to indicate that the allocation amount is 30 blocks, the syntax is:

\$XABALL ALQ=30

\$XABALL AOP

6.5.4 Allocation Option

The AOP parameter sets the allocation option field, which lets you specify a particular type of allocation.

Format

AOP=<CBT,CTG,HRD,ONC>

The AOP parameter can indicate any number of options. When only one option is chosen, angle brackets (< and >) are not required; otherwise, they are required syntax. The allocation options may be specified in any order.

CBT

Contiguous best try; indicates that VAX-11 RMS is to perform the initial allocation (or a later extension) using contiguous blocks, on a "best effort" basis. This overrides the CBT bit in the file processing options (FOP) field of the FAB.

THE EXTENDED ATTRIBUTE BLOCKS

CTG

Contiguous; indicates that the initial allocation (or later extension) must use contiguous blocks only; the allocation fails if the requested number of contiguous blocks is not available. If this is the initial allocation, the file is marked contiguous. Overrides the CTG bit in the file-processing options field of the FAB.

HRD

Hard; indicates that if the requested alignment cannot be performed, an error will be returned. The default is that allocation is to be performed as near as possible to the requested alignment.

NOTE

The HRD option is applicable only to CYL and LBN alignment boundary types, specified by the ALN parameter of the allocation XAB.

ONC

On cylinder boundary; indicates that VAX-11 RMS is to start the allocation on any available cylinder boundary.

For example, suppose you want 30 blocks allocated contiguously starting at logical block number 1024 with an error returned if not possible. You would write:

```
$XABALL  ALQ=30,          -; allocation amt
          ALN=LBN         -; start at logical blk. no.
          LOC=1024        -; 1024
          AOP=<CTG,HRD>   ; contig. or rtn. error
```

Each allocation request option has its own symbolic bit offset and mask value.

\$XABALL BKZ

6.5.5 Bucket Size

The BKZ parameter initializes the bucket size field, which is used only with the relative and indexed file organizations. When you create a relative or indexed file, you specify the bucket size field before issuing the \$CREATE macro instruction. For a relative file, the BKZ parameter specifies the bucket size because a relative file may have only one area. However, for an indexed file, the BKZ parameter specifies the bucket size for the area described by the allocation XAB; this allows you to vary the size of buckets among the multiple areas of your indexed file. When you open an existing file, VAX-11 RMS sets this field to the defined size of the buckets in the file for a relative file or the defined size of the buckets in this area (defined by the AID parameter) for an indexed file.

The value in this field overrides the contents of the bucket size field (BKS) of the FAB on a Create service (see Section 4.2.3).

THE EXTENDED ATTRIBUTE BLOCKS

Format

BKZ=bucket-size

bucket-size

A numeric value, in the range of 0 through 32, representing the number of blocks in a bucket. If this parameter is omitted or if a value of 0 is used, then a default size will be used equal to the minimum number of blocks required to contain a single record.

For example, to specify a bucket size of two blocks, you would write:

```
$XABALL BKZ=2
```

\$XABALL DEQ

6.5.6 Default Extension Quantity

The DEQ parameter initializes the default extension quantity field, which specifies the number of blocks to add to the file whenever it is extended automatically.

The value in this field overrides the contents of the default extension quantity field (DEQ) of the FAB (see Section 4.2.6).

Format

DEQ=extension-quantity

extension-quantity

The number of blocks to be added when automatic extension is required. This number must be in the range of 0 through 65,535. If you specify 0, the file will be extended using a VAX-11 RMS-determined default extension value.

For example, to specify a default extension quantity of 50 blocks, you would write:

```
$XABALL DEQ=50
```

\$XABALL LOC

6.5.7 Location

The LOC parameter initializes the location field, indicating the starting point for file allocation. The exact interpretation of this field depends on the contents of the alignment boundary type field (ALN) (see Section 6.5.3). VAX-11 RMS uses the contents of the location field on a \$CREATE or \$EXTEND macro instruction, but only if the alignment boundary type field (ALN) is also initialized.

THE EXTENDED ATTRIBUTE BLOCKS

Format

LOC=number

number

The starting point for the allocation is determined from the contents of the alignment boundary type field as follows:

- If CYL is specified for the ALN parameter, the LOC number specified is the starting cylinder number where the allocation is to start, in the range of 0 through the maximum cylinder number on the volume.
- If LBN is specified for the ALN parameter, the LOC number specified is the logical block number where the allocation is to start, in the range of 0 through the maximum number of blocks on the volume.
- If VBN or RFI is specified for the ALN parameter, the LOC number specified is the virtual block number where the allocation is to start, in a range from 1 through the maximum number of blocks in the file. This is used only in conjunction with a \$EXTEND macro instruction. If the number 0 is specified, or if the number is omitted during an extend operation, VAX-11 RMS extends as near to the end of the file as possible.

For example, to indicate that you want to allocate 30 contiguous blocks starting at or near logical block 1024, you would write:

```
$XABALL ALQ=30           -; allocate 30 blocks
        ALN=LBN         -; start at logical block
        LOC=1024        -; number 1024
        AOP=CTG         ; contiguously
```

6.5.8 Relative File Identifier

The RFI parameter sets the related file identification field, which lets you allocate files close to other files.

Format

RFI = <F1, F2, F3>

F1,F2,F3

The three-word file identification value of the related file. For a discussion of file identification value, see Section 7.3. A value of 0,0,0 (the default) indicates that the current file is to be used. Specifying RFI in the ALN field and specifying RFI=<0,0,0> is equivalent to specifying ALN=VBN.

The angle brackets are related syntax. This file is created or extended as near to the specified related file as possible at the virtual block number specified by the LOC parameter.

The RFI is ignored if the ALN parameter is not set to RFI.

\$XABALL VOL

6.5.9 Relative Volume Number

The VOL parameter initializes the relative volume number field. It indicates the specific member of a volume set upon which the file is to be allocated.

Format

VOL=volume-number

volume-number

An integer in the range 0 through 65535. Assembly-time default is 0, specifying the "current" member of the volume set.

For example, to indicate that the file is to reside on relative volume number 3 of the volume set, you would write:

```
$XABALL VOL=3,ALQ=30,ALN=CYL,LOC=1
```

NOTE

Volume placement will be performed only if an alignment type (in the ALN field) is also specified. If the ALN field is zero, placement of the file within the volume set is at the discretion of the system, regardless of the contents of the VOL field.

\$XABKEY

6.6 KEY DEFINITION XAB

The \$XABKEY macro instruction allocates and initializes an XAB that defines the key fields of an indexed file at file creation; it also allows retrieval of the key definition at file open and display. Each key definition XAB describes one key of an indexed file.

When you create an indexed file, you must set the contents of the fields of this XAB before you issue the \$CREATE macro instruction. Further, you must provide one key definition for each key that you want the file to have. Since every indexed file must have at least one key, the primary key, you will always require at least one key definition XAB.

When you open an existing indexed file or issue a Display operation for such a file, you use key definition XABs only if you want VAX-11 RMS to provide your program with one or more of the key definitions specified when the file was created.

Table 6-5 summarizes the fields that comprise the key definition XAB.

THE EXTENDED ATTRIBUTE BLOCKS

Table 6-5
Key Definition Extended Attribute Block Fields

Field Name	Field Size	Description	Offset
DAN	byte	Data bucket area number	XAB\$B_DAN
DBS ¹	byte	Data bucket size	XAB\$B_DBS
DFL	word	Data bucket fill size	XAB\$W_DFL
DTP	byte	Data type of the key	XAB\$B_DTP
DVB ¹	longword	First data bucket start virtual block number	XAB\$L_DVB
FLG	byte	Key options flag	XAB\$B_FLG
IAN	byte	Index buckets area number	XAB\$B_IAN
IBS ¹	byte	Index bucket size	XAB\$B_IBS
IFL	word	Index bucket file size	XAB\$W_IFL
KNM	longword	Key name buffer address	XAB\$L_KNM
LAN	byte	Lowest level of index area number	XAB\$B_LAN
LVL ¹	byte	Level of root buckets	XAB\$B_LVL
MRL ¹	word	Minimum record length	XAB\$W_MRL
NSG ¹	byte	Number of key segments	XAB\$B_NSNG
NUL	byte	Null key value	XAB\$B_NUL
POS	word	Key position	XAB\$W_POS0 through XAB\$W_POS7
REF	byte	Key of reference	XAB\$B_REF
RVB ¹	longword	Root bucket start virtual block number	XAB\$L_RVB
SIZ	byte	Key size	XAB\$B_SIZ0 through XAB\$B_SIZ7
TKS ¹	byte	Total key field size	XAB\$B_TKS

¹Indicates nonuser-initialized field

THE EXTENDED ATTRIBUTE BLOCKS

Format

OPERATION	PARAMETERS
label:\$XABKEY	DAN=area-number DFL=bytes DTP=data-type-code FLG=<CHG, DUP, NUL> IAN=area-number IFL=bytes KNM=address LAN=area-number NUL=value POS=<position, ...> REF=value SIZ=<size, ...>

\$XABKEY DAN

6.6.1 Data Bucket Area Number

The DAN parameter initializes the data bucket area number field of the key definition XAB. You use this parameter to specify the area of the file that the data buckets are to reside in only when both of the following are true:

- You are creating a new indexed file
- You are using allocation XABs (described in Section 6.5) to define areas

When the key definition XAB describes the primary key, the data level of the index consists of buckets that contain the actual data records of the file. However, when the key definition describes an alternate key, the data level of the index consists of buckets in which VAX-11 RMS maintains pointers to the actual data records.

Format

DAN=area-number

area-number

A numeric value in the range 0 through 254, representing an identification number contained in the AID field of an allocation XAB present in the same chain (see Section 6.5.2). The default is 0, that is, area 0.

THE EXTENDED ATTRIBUTE BLOCKS

For example, to indicate that these data buckets are to reside in area 3 of an indexed file, you would write:

```
$XABKEY DAN=3
```

\$XABKEY DFL

6.6.2 Data Bucket Fill Size

The DFL parameter initializes the data bucket fill size field of the key definition XAB. When you create an indexed file, you use this parameter to specify the number of bytes (of data) you want in each data level bucket. If you specify less than the total possible bucket size, you thereby indicate that the data buckets are to contain some amount of free space. At run time, VAX-11 RMS follows the fill size specified at Create time only if the RAB\$V_LOA bit is set in the record processing options (ROP) field of the RAB. The ROP field is described in Section 5.

When the key definition XAB describes the primary key, the DFL field describes the space in the buckets containing actual user data records. When the key definition XAB describes an alternate key, the DFL field describes the space in the buckets containing pointers to the user data records.

It is advantageous to use the DFL field in the following situation:

If you expect to perform numerous Put and Update operations on the file after it has been initially populated, you can minimize the resultant movement of records (known as bucket splitting) by specifying less than the maximum bucket fill size at Create time. To utilize the free space thereby reserved in the buckets, programs that perform Put or Update operations on the file should not place the value RAB\$V_LOA in the ROP field of the RAB.

Format

```
DFL=bytes
```

bytes

A numeric value representing the maximum number of bytes (of data) in a data bucket. The maximum possible fill size is the bucket size, in blocks, multiplied by 512. The assembly-time default value is 0, which is interpreted by VAX-11 RMS as meaning the maximum available space (i.e., no unused space). If the specified size is not zero, but is less than one half of the bucket size (in bytes), then the fill size used will be one half of the bucket size.

For example, to specify that each bucket at the data level is to be filled to a maximum of 400 bytes, you would write:

```
$XABKEY DFL=400
```

\$XABKEY DTP**6.6.3 Key Data Type**

The DTP parameter initializes the data type of the key field of the XAB. When you create an indexed file, you use this parameter to specify the type of data in the record key field.

Key field data types and the data type codes are summarized and the associated global symbols are listed in Table 6-6.

Table 6-6
Key Field Data Types, Data Type Codes and Global Symbols

Key Field Data Type	Data Type Code	Global Symbol
String	STG	XAB\$C_STG
Signed 2-byte integer	IN2	XAB\$C_IN2
Signed 4-byte integer	IN4	XAB\$C_IN4
Unsigned 2-byte binary	BN2	XAB\$C_BN2
Unsigned 4-byte binary	BN4	XAB\$C_BN4
Packed decimal	PAC	XAB\$C_PAC

String data type (STG) is defined as a left-justified string of unsigned 8-bit bytes.

The string key field consists of from one through eight disjointed key field segments (see Sections 6.6.11 and 6.6.13).

Integer, binary, and packed decimal key fields must be a contiguous set of bytes.

The null value (that is, NUL option in FLG parameter is set) for integer, binary, and packed decimal is zero and the NUL parameter (field) is ignored (see Sections 6.6.5 and 6.6.10).

A packed decimal is a contiguous sequence of bytes and is specified by two attributes: the address, A, of the first byte of the string and a length, L, that is the number of digits in the packed decimal. The bytes of a packed decimal are divided into two 4-bit fields that must contain decimal digits, except for the first four bits (0 through 3) of the last (highest addressed) byte, which must contain a sign. The representation for the digits and signs is shown in Table 6-7.

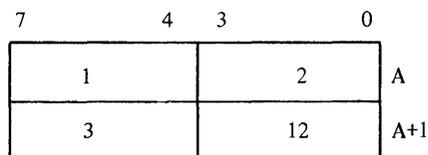
THE EXTENDED ATTRIBUTE BLOCKS

Table 6-7
Packed Decimal Digits and Signs Representation

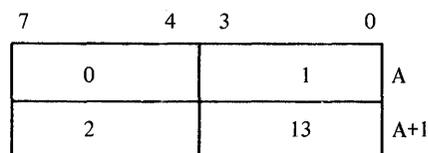
Digit or Sign	Decimal	Hex
0	0	0
1	1	1
2	2	2
3	3	3
4	4	4
5	5	5
6	6	6
7	7	7
8	8	8
9	9	9
+	10, 12, 14 or 15	A, C, E or F
-	11 or 13	B or D

The preferred sign representation is 12 for plus (+) and 13 for minus (-). The length L is the number of digits in the packed decimal string (not counting the sign) and must be in the range 0 through 31. When the number of digits is even, it is required that an extra 0 digit appear in the last four bits (4 through 7) of the first byte. Again the length in bytes of the packed decimal is $L/2 + 1$. The value of a 0-lengthpacked decimal is 0; it contains only the sign byte, which also includes the extra 0 digit.

The address, A, of the packed decimal specifies the byte containing the most significant digit in its high nibble. Digits of decreasing significance are assigned to increasing byte addresses and from high to low within a byte. Thus +123 has length 3 and is represented as follows:



and -12 has length 2 and is represented as follows:



Integer and binary key field data have the following formats:

- IN2: LSB at A
 MSB and sign at A+1
- IN4: LSB at A
 MSB and sign at A+3
- BN2: LSB at A
 MSB at A+1
- BN4: LSB at A
 MSB at A+3

THE EXTENDED ATTRIBUTE BLOCKS

Format

DTP=data-type-code

data-type-code

One of the following, as appropriate:

STG, string (left-justified, unsigned 8-bit bytes), this is the default
IN2, signed 2-byte integer key data
IN4, signed 4-byte integer key data
BN2, unsigned 2-byte binary key data
BN4, unsigned 4-byte binary key data
PAC, packed decimal key data

For example, to specify that the key data type is a signed 4-byte integer, you would write:

```
$XABKEY DTP=IN4
```

\$XABKEY FLG

6.6.4 Key Options Flag

The FLG parameter initializes the key options flag field of the key definition XAB. When you create an indexed file, you specify the following optional characteristics of the key represented by this XAB:

- Key values can change
- Duplicate key values are permitted
- Null key value

Format

FLG CHG, DUP, NUL

Option

One of the following as appropriate:

CHG

The key value within the record in the file can be changed by a program during a \$UPDATE operation. This option can be specified only for alternate keys.

DUP

The key value within the record in the file may have the same key value as another record (or other records) within the file.

THE EXTENDED ATTRIBUTE BLOCKS

NUL

The NUL field of the XAB contains a null key value if the key data type is string. If the key data type is other than string (i.e., integer, binary, or packed decimal), then the null key value is 0. This option can be specified only for alternate keys. Refer to Section 6.6.10 for a description of the XAB NUL field.

The allowed combinations of the changeable key values and duplicate key values options depend on the type of key (that is, primary or alternate) represented by this XAB; Table 6-8 summarizes these combinations.

Table 6-8
Key Options Flag Combinations

Key Type	Combinations			
	CHG + DUP	CHG + NO DUP	NO CHG + DUP	NO CHG + NO DUP
Primary	Error	Error	Allowed	Default
Alternate	Default	Allowed	Allowed	Allowed

The assembly-time defaults for the FLG field depend on the key of reference specified by the REF field.

The defaults for a primary key are as follows:

- Duplicate key values are not allowed
- Key values cannot change

The defaults for an alternate key are as follows:

- Duplicate key values are allowed
- Key values can change
- No null key values

The defaults are applied only if the entire FLG field is defaulted. Consider the following:

```
KEY 1: $XABKEY REF=1 POS=0 SIZ=10
```

This specifies the key for alternate index 1, and therefore the macro will default the FLG field to allow duplicates and changes. However, if a FLG option is explicitly referenced, the results are different:

```
KEY 1: $XABKEY REF=1 POS=0 SIZ=10 FLG=CHG
```

In this case, only the CHG option will be set in the FLG field, leaving the DUP option cleared. This means that duplicates will not be allowed on this key.

When you specify more than one option with the FLG parameter, you must enclose the options in angle brackets. The options can be specified

THE EXTENDED ATTRIBUTE BLOCKS

in any order. When you specify only one option, no angle brackets are required.

For example, to specify that duplicate key values are allowed, that a null key value is allowed, and that key values cannot change (through absence of CHG), you would write the following:

```
$XABKEY FLG=<DUP,NUL>
```

Each key option flag operation has its own symbolic bit offset and mask value.

Special Note

VAX-11 RMS will allow alternate indices which do not permit duplicate key values in that index, but do permit key values to change on \$UPDATE operations. RMS-11 (as opposed to VAX-11 RMS) does not allow this particular combination of attributes for alternate indices. This factor should be considered when creating files with VAX-11 RMS which may also be processed by RMS-11.

\$XABKEY IAN

6.6.5 Index Bucket Area Number

The IAN parameter initializes the index bucket area number field of the key definition XAB. When you create an indexed file, you use this parameter to specify the area of the file that the index buckets are to reside in only when both of the following are true:

- You are creating a new indexed file.
- You are using allocation XABs (described in Section 6.5) to define areas.

When the key definition XAB describes the primary key, the index level of the index consists of all levels of the tree- (pyramid-)structured primary index down to and including the level containing pointers to the user data records themselves. However, when the key definition describes an alternate key, the index level of the index comprises all levels of the pyramid-structured alternate index down to, but not including, the level containing buckets in which VAX-11 RMS maintain pointer arrays describing the user data records. Refer to the LAN parameter for a description of how to place the lowest level of the index in a location separate from the higher levels.

Format

```
IAN=area-number
```

area-number

A numeric value in the range 0 through 254, representing an identification number contained in the AID field of an allocation XAB present in the same chain (see Section 6.5.2). The default is 0, that is, area 0.

THE EXTENDED ATTRIBUTE BLOCKS

For example, to indicate that these index buckets are to reside in area 3 of an indexed file, you would write:

```
$XABKEY IAN=3
```

\$XABKEY IFL

6.6.6 Index Bucket Fill Size

The IFL parameter initializes the index bucket fill size field of the key definition XAB. When you create an indexed file, you use this parameter to specify the number of bytes you want in each index bucket. If you specify less than the total possible bucket size, you indicate that the index buckets are to contain some amount of free space. At run time, VAX-11 RMS adheres to the fill size specified at \$CREATE time only if the RAB\$V_LOA bit is set in the record-processing options (ROP) field of the RAB. The ROP field is described in Chapter 5.

When the key definition XAB describes the primary key, the IFL field describes the space in the buckets in all levels of the primary index down to and including the level containing pointers to the user data records. When the key definition XAB describes an alternate key, the IFL field describes the space in the buckets in all levels of the alternate index down to, but not including, the level containing buckets in which VAX-11 RMS maintains pointer arrays describing the user data records.

It is advantageous to use the IFL field in the following situation:

If you expect to perform numerous \$PUT and \$UPDATE operations on the file after it has been initially populated, you can minimize the resultant movement of records (known as bucket splitting) by specifying less than the maximum bucket fill size at \$CREATE time. To utilize the free space thereby reserved in the buckets, programs that perform \$PUT or \$UPDATE operations on the file should not place the value RAB\$V_LOA in the ROP field of the RAB.

Format

```
IFL=bytes
```

bytes

A numeric value representing the maximum number of bytes in an index bucket. The maximum possible fill size is the bucket size, in blocks, multiplied by 512. The default value is 0, which is interpreted by VAX-11 RMS as meaning the maximum available space (that is, no unused space). If the specified size is not zero, but is less than one half of the bucket size (in bytes), then the fill size used will be one half of the bucket size.

For example, to specify that each index bucket is to be filled to a maximum of 256 bytes, you would write:

```
$XABKEY IFL=256
```

\$XABKEY KNM

6.6.7 Key Name Address

The KNM parameter initializes the key name buffer address field of the key definition XAB. When you define a key during creation of an indexed file, you can associate any 32-character string you choose with the key field represented by the XAB. VAX-11 RMS never examines this character string, but it retains it in the file as part of the key definition information.

Format

KNM=address

address

The symbolic address of a buffer, which must always be at least 32 bytes long. A value of 0 in this field indicates that no key name is defined during a \$CREATE operation or is to be displayed during a \$OPEN or \$DISPLAY operation.

For example, if the key buffer area has a label of KEYBUF, you would write:

\$XABKEY KNM=KEYBUF

\$XABKEY LAN

6.6.8 Lowest Level of Index Area Number

The LAN parameter initializes the lowest level of index area number field of the key definition XAB. It permits you to separate the lowest level (level 1) of the index from all higher levels (levels 2 +) of the index in an indexed file; that is, you can use the LAN parameter to specify an area of the index wherein the lowest level of the index will reside, separate from the area (or areas) specified by the IAN parameter (wherein higher levels of the index will reside). The IAN parameter is described in Section 6.6.6.

You can utilize the LAN parameter only when both of the following are true:

- You are creating a new indexed file.
- You are using allocation XABs (described in Section 6.5) to define areas.

NOTE

The bucket size of the area specified by the LAN parameter must be the same as the bucket size specified by the IAN parameter.

THE EXTENDED ATTRIBUTE BLOCKS

Format

LAN=area-number

area-number

A numeric value in the range 0 through 254, representing an identification number contained in the AID field of an allocation XAB present in the same chain (see Section 6.5.2). The assembly-time default is 0; that is, the lowest level of the index will occupy the same area of the file as the remainder of the index.

For example, to indicate that the lowest level of the index is to reside in area 3 of an indexed file, you would write:

```
$XABALL AID=3      ;area identification
$XABKEY IAN=5      -;index area number
          LAN=3      ;lowest level of index area number
```

\$XABKEY NUL

6.6.9 Null Key Value

The NUL parameter initializes the null field of the key definition XAB. Normally, VAX-11 RMS updates all indexes to reflect the values in the corresponding key fields of the records written to an indexed file. The NUL parameter, however, allows you to instruct VAX-11 RMS not to make an entry in an alternate index if a record being entered in an indexed file contains the specific (null) alternate key value. The following prerequisites must be satisfied for you to use the NUL parameter:

- The XAB must define an alternate key.
- The NUL option of the FLG parameter must have been set at file creation (refer to Section 6.6.5 for a description of the FLG parameters).
- The key data type must be string.

If the above conditions are met, alternate index entries will not be made for those alternate key values in which every byte of the key matches the null key value for that index. Non-string key data types use 0 for the null key value.

Format

NUL=value

value

Any user-selected character value

For example, to indicate that a record with an alternate key value of 32 (ASCII blank) is not to have an entry made for it in the associated alternate index (in this case the second alternate index), you would write:

```
$XABKEY FLG=NUL    -;set null flag
          NUL=32    -;null key value
          REF=2     ;second alternate key
```

\$XABKEY POS

6.6.10 Key Position

The POS parameter initializes the key position field of the key definition XAB. The key position field defines the location of the key within each record of an indexed file, and is eight words long. Two types of keys can be defined: simple keys and segmented keys.

A simple key is a single string of contiguous bytes in the records. The first word of the position field specifies the starting position of the string and the remaining words contain 0s. You can use simple keys with any data type (see Section 6.6.4).

Segmented keys can be used only with key fields that contain string data. A segmented key consists of two through eight strings of bytes in the record. Each individual string (segment) is a set of contiguous bytes, but the strings do not need to be contiguous; additionally, the strings can be in any order and may overlap. Each successive word of the position field specifies a starting position of one of the segments. When processing records that contain segmented keys, VAX-11 RMS regards the key field as a single, logically contiguous string beginning with the first segment and ending with the last.

You should note that the key position and the key size field (see Section 6.6.13) must define an equal quantity of key position values and key size values.

Format

```
POS=position
```

or

```
POS=<POSITION0,POSITION1,...,POSITION7>
```

position

Is a numeric value representing the starting (byte) position of the key within each record. The first byte of a record is represented by the value 0, the second by the value 1, etc. A simple key has only one starting position, while a segmented key may have up to eight starting positions.

For example, to indicate that a record contains a simple key which starts in the first byte of each record, you would write:

```
$XABKEY POS=0, -; key starts in first byte
        SIZ=8   ; key length 8 bytes
```

To indicate that a record contains a segmented key consisting of 4 segments with the first segment starting in the 20th byte, the second segment starting in the 14th byte, the third segment starting in the first byte, and the fourth segment starting in the 29th byte, you would write:

```
$XABKEY POS=<19,13,0,28>, -; segmented key
        SIZ=<8,2,5,32>   ; length in bytes
```

You must include the angle brackets for multiple argument key positions.

THE EXTENDED ATTRIBUTE BLOCKS

The offsets for these fields are:

XAB\$W_POS0,...,XAB\$W_POS7

\$XABKEY REF

6.6.11 Key of Reference

The REF parameter initializes the key of reference field in the key definition XAB. The key of reference field identifies which key (that is, primary, first alternate, second alternate, and so on) in an indexed file is defined by the XAB. Since REF is a reserved key word in the BLISS language, the BLISS XABKEY macro uses the mnemonic KREF to reference this field.

NOTE

VAX-11 RMS can process an indexed file with 255 defined keys; however, you should be aware that each key field defined has associated with it a cost in processing and I/O time. The time to build and maintain the index for the key field and the disk storage required to contain the index for each key field should be considered when you decide whether the field should be an alternate key field. A file with six to eight defined keys (the primary and five to seven alternate keys) should be considered as a maximum; a file with two or three defined keys is normal.

Format

REF=value

value

A numeric value in the range 0 through 254 indicating which key is represented by the XAB. A value of 0 indicates the primary key, 1 indicates the first alternate key, 2 indicates the second alternate key, and so on. For the \$CREATE and \$EXTEND macro instructions, the key references must be listed consecutively in ascending order. The order is irrelevant for the \$OPEN and \$DISPLAY macro instructions.

For example, to indicate the primary key, you would write:

\$XABKEY REF=0

\$XABKEY SIZ

6.6.12 Key Size

The SIZ parameter initializes the key size field of the key definition XAB. The key size field defines the length (in bytes) of the key (whose starting position is defined in the key position field of the same XAB) within each record of an indexed file. Two types of keys can be defined: simple keys and segmented keys (see Section 6.6.11). The key size field defining a simple key will contain only one key size value. The key size field defining a segmented key must contain a key size value for each segment of the key. You should note that the key size field and the key position field (see Section 6.6.11) must contain an equal quantity of key size values and key position values. VAX-11 RMS associates the first key position value specified with the first key size value specified which together define the location and length of the first segment of a segmented key, and so forth.

Format

```
SIZ=size0
```

or

```
SIZ=<SIZE0,SIZE1,...,SIZE7>
```

size

A numeric value representing the length, in bytes, of the key within the record. Up to eight values can be assigned.

When the data type of the key (see Section 6.6.4) is string, the total size (sum of SIZE,SIZE,...) of the key must be less than 256 bytes.

When the data type of the key (see Section 6.6.4) is 2-byte integer or 2-byte binary, size0 must equal 2 and size1 through size7 must be 0s. If size0 is 0, it is defaulted to 2.

When the data type of the key (see Section 6.6.4) is 4-byte integer or 4-byte binary, size0 must equal 4 and size1 through size7 must be 0s. If size0 is 0, it is defaulted to 4.

When the data type of the key (see Section 6.6.4) is packed decimal, the size specified by size0 must be from 1 through 16 and size1 through size 7 must be 0s.

For example, to indicate that a record contains a simple key eight bytes long, you would write:

```
$XABKEY POS=0, -; key starts in first byte
        SIZ=8   ; key length 8 bytes
```

To indicate that a record contains a segmented key consisting of 4 segments with the first segment 8 bytes long, the second segment 2 bytes long, the third segment 5 bytes long, and the fourth segment 32 bytes long, you would write:

```
$XABKEY POS=<19,13,0,28>, -; KEY SEGMENT START LOCATIONS
        SIZ=<18,2,5,32>   ; KEY LENGTH IN BYTES
```

THE EXTENDED ATTRIBUTE BLOCKS

The offsets for these fields are:

XAB\$B_SIZ0, ..., XAB\$B_SIZ7

6.6.13 Noninitializable Key Fields

The following list describes the fields that are output fields only. VAX-11 RMS sets them for you, when you perform a DISPLAY or an OPEN operation.

DBS

Data bucket size field. When a key definition XAB is present during an open or display operation, VAX-11 RMS sets this field to the size of the data level (level 0) buckets, in virtual blocks, for the key described by the XAB.

DVB

First data bucket start virtual block number. When a key definition XAB is present during an open or display operation, VAX-11 RMS sets this field to the start virtual block number for the first data level bucket for the key described by the XAB.

IBS

Index bucket size. When a key definition XAB is present during an open or display operation, VAX-11 RMS sets this field to the size of the index level (level 1 to n) buckets, in virtual blocks, for the key described by the XAB.

LVL

Level of root bucket. When a key definition XAB is present during an open or display operation, VAX-11 RMS sets this field to the level of the root bucket for the key described by the XAB.

MRL

Minimum record length. When a key definition XAB is present during an open or display operation, VAX-11 RMS sets this field to the minimum record length in bytes, which will totally contain the key field for the key described by the XAB.

If the key described by the XAB is the primary key (REF=0), then a record must be equal to or greater than the minimum record length returned in MRL to be inserted/updated in the file.

If the key described by the XAB is an alternate key (REF=1 to n), then a record must be equal to or greater than the minimum record length returned in MRL to be recorded in the associated index for that alternate key.

NSG

Number of key segments. When a key definition XAB is present during an open or display operation, VAX-11 RMS sets this field to the number of key segments that make up the key field for the key described by the XAB (see Section 6.6.11). This field must not be altered.

THE EXTENDED ATTRIBUTE BLOCKS

RVB

Root index bucket start virtual block number. When a key definition XAB is present during an open or display operation, VAX-11 RMS sets this field to the start virtual block number for the root bucket of the index for the key described by the XAB.

TKS

Total key size. When a key definition XAB is present during an open or display operation, VAX-11 RMS sets this field to the total key size, in bytes (the sum of SIZ0 through SIZ7), for the key described by the XAB (see Section 6.6.13).

\$XABSUM

6.7 SUMMARY XAB

The \$XABSUM macro instruction allows you to determine the number of keys and/or the number of allocation areas defined and the prologue version number for an existing file.

The summary XAB is ignored with a \$CREATE macro call. However, one summary XAB can be associated with a FAB at the time a \$OPEN or \$DISPLAY macro call is issued. The presence of this XAB during these calls allows VAX-11 RMS to return to your program the total number of keys and allocation areas defined and the prologue version number when the file was created.

Format

```
OPERATION          PARAMETERS
label:$XABSUM      NXT=address
```

Table 6-9 summarizes the fields in the summary XAB.

NOTE

The summary XAB is used only with indexed files.

Table 6-9
Summary Extended Attribute Block Fields

Field Name	Field Size	Description	Offset
BLN	byte	Block length	XAB\$B_BLN
COD	byte	Type code	XAB\$B_COD
NOA	byte	Number of allocation areas defined for the file	XAB\$B_NOA
NOK	byte	Numbers of keys defined for the file	XAB\$B_NOK
NXT	longword	Next XAB address	XAB\$L_NXT
PVN	word	Prologue version number	XAB\$W_PVN

\$XABFHC

6.8 FILE HEADER CHARACTERISTICS XAB

The \$XABFHC macro instruction allocates and initializes a file header characteristics XAB. You can use this block to display information about the file as stored in the file header.

VAX-11 RMS copies the file characteristics into this XAB whenever an operation is performed with a \$OPEN or \$DISPLAY macro instruction. The field is then available for you to examine during processing. Note that for shared sequential files, the values in the end-of-file block, first free byte in the end-of-file block, and longest record length fields correspond to the values at the time of the last close or flush service.

On a Create service, only the longest record length field of this XAB is used as an input attribute, and then only if the record format is not fixed length.

Format

OPERATION	PARAMETERS
label: \$XABFHC	NXT=address

Table 6-10 summarizes the fields in the file header characteristics XAB. Note that many of these fields are also available in the FAB.

THE EXTENDED ATTRIBUTE BLOCKS

Table 6-10
File Header Characteristics
Extended Attribute Block Fields

Field Name	Field Size	Description	Offset
ATR	byte	Record attributes; equivalent to the RAT field of the FAB	XAB\$B_ATR
BKZ	byte	Bucket size; equivalent to the BKS field of the FAB	XAB\$B_BKZ
BLN ²	byte	Block length	XAB\$B_BLN
COD ²	byte	Type code	XAB\$B_COD
DXQ	word	Default file extension quantity; equivalent to the DEQ field of the FAB	XAB\$W_DXQ
EBK	longword	End-of-file block	XAB\$L_EBK
FFB	word	First free byte in the end-of-file block	XAB\$W_FFB
HBK	longword	Highest virtual block in the file; the execution of a \$OPEN macro instruction sets the allocation quantity field of the FAB to this value	XAB\$L_HBK
HSZ	byte	Fixed length control header size; equivalent to the FSZ field of the FAB	XAB\$B_HSZ
LRL	word	Longest record length	XAB\$W_LRL
MRZ	word	Maximum record size; equivalent to the MRS field of the FAB	XAB\$W_MRZ
NXT ¹	longword	Next XAB address	XAB\$L_NXT
RFO	byte	File organization and record format; combines the RFM and ORG fields of the FAB	XAB\$B_RFO
SBN	longword	Starting logical block number for the file if it is contiguous, otherwise this field is 0	XAB\$L_SBN

¹This field can be initialized at assembly time.

²Indicates that this field is set automatically by the type of macro instruction.

\$XABRDT**6.9 REVISION DATE AND TIME XAB**

The \$XABRDT macro instruction allocates and initializes an XAB for revision date and time. This XAB operates much like the date and time XAB (see Section 6.3) when input to the \$OPEN, \$DISPLAY, or \$CREATE macro instructions. However, when you gain access to a file for writing, issuing a \$CLOSE macro instruction for that file causes the revision date and time to be set from the current date and time and the revision number to be incremented. Thus, any revision date and time you specify through the XAB on a \$CREATE macro instruction is lost.

For this reason, you can input the revision date and time XAB to the \$CLOSE macro instruction and cause the file's revision date and time and revision number to take on the specified values.

Table 6-11 summarizes the fields in the revision date and time XAB.

Table 6-11
Revision Date and Time Extended Attribute Block Fields

Field Name	Field Size	Description	Offset
BLN ²	byte	Block length	XAB\$B_BLN
COD ²	byte	Type code	XAB\$B_COD
NXT	longword	Next XAB address	XAB\$L_NXT
RDT ¹	quadword	Revision date and time	XAB\$Q_RDT
RVN ¹	word	Revision number	XAB\$W_RVN

¹ Indicates no assembly time initialization.

² Indicates that this field is set automatically by the type of macro instruction.

Format

OPERATION	PARAMETERS
label: \$XABRDT	NXT=address

THE EXTENDED ATTRIBUTE BLOCKS

6.9.1 Revision Date and Time

VAX-11 RMS sets certain values for the revision date and time, and returns them in the revision date and time XAB fields for your inspection. You can override these system-supplied values through the use of a revision date and time XAB as input to a \$CLOSE or \$CREATE macro instruction. However, the \$XABRDT macro instruction does not contain parameters for the assembly-time initialization of XABRDT fields. As outlined in Table 6-8, these fields are:

- Revision date and time (RDT) -- this is a 64-bit binary field, indicating the date and time at which the file was last updated.
- Revision Number (RVN) -- this field provides the number of times this file was opened for write operations.

The following tables describe how the fields of the XABRDT are used by the RMS file processing macro instructions.

OPERATION	RDT	RVN
CLOSE	INPUT	INPUT
CREATE	INPUT	INPUT
DISPLAY	OUTPUT	OUTPUT
ERASE	not used	not used
EXTEND	not used	not used
OPEN	OUTPUT	OUTPUT

If you specify a revision date and time of zero, VAX-11 RMS will substitute the current date and time on the close operation. If, however, the XABRDT existed on an open (or a display) operation, the RDT field will have been filled with the existing file's RDT value. Since the RDT has been filled (and is no longer zero), VAX-11 RMS will not substitute the current date and time.

The XABRDT should be used only in those cases in which you want to specify a new, (and non-default) RDT and RVN. These fields may be set anytime after the open operation, but must be set before the close operation. If the fields are set before the open operation, the files' existing RDT and RVN will override the values you specified. If you only want to examine the contents of the RDT and RVN fields, you can do so by simply consulting the information already contained in the appropriate fields of the XABDAT block.

CHAPTER 7

THE NAME BLOCK

This chapter describes the Name (NAM) Block, its fields, and the macro instruction and parameters that initialize the fields at assembly time.

7.1 THE PURPOSE OF THE NAME BLOCK

The NAM block contains supplementary information for use with the file specification, and is useful as a means to facilitate file opening. The fields of the NAM block include the following information:

- Device identification
- Directory identification
- File identification
- Expanded and resultant file name strings
- Address of a related file's NAM block
- Wild card character context

To use a NAM block, you must specify its symbolic address as the value in the name block address field (NAM parameter) of the associated FAB.

The \$NAM macro instruction allocates a NAM block. At assembly time, you can initialize the fields in the NAM block through keyword parameters. For run-time access to these fields, you can use the keyword parameters with the \$NAM_STORE macro instruction (see Chapter 14), or the symbolic offsets.

Table 7-1 summarizes the fields in the NAM block. Some of these fields, however, are set by VAX-11 RMS or are static; therefore, you cannot initialize them at assembly time by keyword parameters.

THE NAME BLOCK

Table 7-1
Name Block Fields

Field Name	Field Size	Description	Offset
BID ¹	byte	Block identifier	NAM\$B_BID
BLN ¹	byte	Block length	NAM\$B_BLN
DID ¹	3 words	Directory identification	NAM\$W_DID
DVI ¹	16 bytes	Device identification	NAM\$T_DVI
ESA	longword	Expanded string area address	NAM\$L_ESA
ESL ¹	byte	Expanded string length	NAM\$B_ESL
ESS	byte	Expanded string area size	NAM\$B_ESS
FID ¹	3 words	File identification	NAM\$W_FID
FNB ¹	longword	File name status bits	NAM\$L_FNB
RLF	longword	Related file NAM block address	NAM\$L_RLF
RSA	longword	Resultant string area address	NAM\$L_RSA
RSL ¹	byte	Resultant string length	NAM\$B_RSL
RSS	byte	Resultant string area size	NAM\$B_RSS
WCC ¹	longword	Wildcard context	NAM\$L_WCC

¹Indicates nonuser-initialized field

\$NAM

7.2 NAM BLOCK ALLOCATION

The \$NAM macro instruction allocates and initializes storage for a NAM block. You cannot use this macro instruction within a sequence of executable instructions.

THE NAME BLOCK

Format

OPERATION	PARAMETERS
label: \$NAM	ESA=address ESS=size RLF=nam-address RSA=address RSS=size

label: \$NAM

7.2.1 Label

The label for the \$NAM macro instruction assigns a name for a particular NAM block and thus provides a symbolic address to be stored in the name block address field of the FAB.

For example, if a label of NMBLK is used for a NAM block, the syntax is:

```
$FAB MRS=512,MRN=1000,NAM=NMBLK,ORG=REL
```

A label must be separated from the \$NAM macro name by a colon (:).

\$NAM ESA

7.2.2 Expanded String Area Address

The ESA parameter initializes the expanded string area address field of the NAM block, which contains the symbolic address of a user-allocated buffer. This buffer receives the file specification string resulting from the translation of logical names and the application of default file specification information to the original file string (file specification string of the FAB). The default file specification information consists of the default file specification string of the FAB, the related file resultant specification string, and the process defaults.

You must specify this field for wild card character processing.

Format

```
ESA=address
```

address

The symbolic address of a buffer in your program to receive the expanded file specification string.

THE NAME BLOCK

For example, if the buffer in your program has a symbolic address of NAMBUF, the syntax is:

```
$NAM ESA=NAMBUF,ESS=32
```

\$NAM ESS

7.2.3 Expanded String Area Size

The ESS parameter initializes the expanded string area size field. This field contains the size of the user-allocated buffer whose address is stored in the expanded string area address field (see Section 7.2.2).

Format

```
ESS=size
```

size

A numeric value representing the size, in bytes, of the user buffer that contains the file specification string, in the range of 0 through 255.

For example, if the user buffer is 32 bytes long, the syntax is:

```
$NAM ESS=32,ESA=NAMBUF
```

The symbolic value NAM\$C MAXRSS defines the maximum possible length of an expanded file specification string.

\$NAM RLF

7.2.4 Related File Nam Block Address

The RLF parameter sets the related file NAM block address field to indicate the address of the NAM block for the related file. This field supports the secondary file concept of the command language (DCL), giving an extra default level in processing file specifications. See Chapter 8 for a description of file specification string parsing.

Format

```
RLF=nam-address
```

nam-address

The symbolic address of the NAM block for the related file.

For example, if the \$NAM macro instruction for the related file NAM block has the label INNAM, the syntax is:

```
$NAM RLF=INNAM
```

\$NAM RSA**7.2.5 Resultant String Area Address**

The RSA parameter sets the resultant string area address field. This field contains the address of a user-allocated buffer that will receive a copy of the resultant file specification string. This string results from the resolution of all system defaults, including version numbers and wild card character substitutions. You must specify this field for wild card processing or when you select the SPL (spool) or SCF (submit) or the DEL (delete on close) options in the FAB.

Format

RSA=address

address

The symbolic address of a buffer in your program that will receive the resultant file specification string.

For example, if the buffer has a label of STRING defining its starting address, the syntax is:

\$NAM RSA=STRING,RSS=48

\$NAM RSS**7.2.6 Resultant String Area Size**

The RSS parameter sets the resultant string area size field. This field defines the length of the user-allocated buffer whose address is contained in the resultant string area address field (see Section 7.2.5).

Format

RSS=size

size

A numeric value representing the size, in bytes, of the buffer that will receive the copy of the file specification string, in the range of 0 through 255.

For example, if the label STRING defines the starting address of a buffer 48 bytes long, the syntax is:

\$NAM RSA=STRING,RSS=48

The symbolic value `NAM$C_MAXRSS` defines the maximum possible length of a resultant file specification string.

THE NAME BLOCK

7.3 NONINITIALIZABLE NAM BLOCK FIELDS

The following list describes the NAM block fields that you cannot initialize at assembly time. Either they are static or VAX-11 RMS sets them for you.

BID

Block identifier field; identifies the block as a NAM block to VAX-11 RMS. The \$NAM macro instruction sets this field to the symbolic value `NAM$C_BID`; you cannot alter this field.

BLN

Block length field; defines the length of the NAM block, in bytes. The \$NAM macro instruction sets this field to the symbolic value `NAM$C_BLN`; you cannot alter this field.

DID

Directory identification field; identifies the directory for the file. VAX-11 RMS outputs this three-word field as part of the \$OPEN, \$CREATE, and \$PARSE macro instructions. If, once you open the file, you want to refer to this directory again, you can do so more quickly by specifying that the NAM block has a valid directory identifier (see Chapter 8).

DVI

Device identification field; defines the device for the file. VAX-11 RMS outputs this field as part of the \$OPEN, \$CREATE, and \$PARSE macro instructions. You can use this field with the file identification field to reopen the file by referring to the NAM block (see Chapter 8). The symbolic value `NAM$C_DVI` gives the length of this field, in bytes.

ESL

Expanded string length field; VAX-11 RMS sets this field as part of the \$OPEN, \$CREATE, and \$PARSE macro instructions. This field is set to the length, in bytes, of the file specification string returned in the buffer whose address is in the expanded string area address field (see Section 7.2.2).

FID

File identification field; provides the identifier of the file. VAX-11 RMS sets this three-word field on a normal open or create operation. You can also set this field before opening the file if you are going to open by file identifier (see Chapter 8).

FNB

File name status bits field; set by VAX-11 RMS to indicate status information about the file as determined by the file specification parsing routine. Each bit within this field denotes a specific status relative to the various components of the file specification. The bits, and the conditions they express, are described in Table 7-2.

Each status bit has its own offset and mask value.

THE NAME BLOCK

RSL

Resultant string length field; VAX-11 RMS sets this field as part of the \$OPEN, \$SEARCH, and \$CREATE macro instructions. This field is set to the length, in bytes, of the file specification string returned in the buffer whose address is in the resultant string area address field (see Section 7.2.5).

WCC

Wild card context field; contains information required for using wild card characters in place of the various file specification components. In particular, this field restarts a directory search to find the next matching file name, type, and/or version number.

THE NAME BLOCK

Table 7-2
File Name Status Bits

Bit Names	Description
DIR_LVL5	Number of sub-directory levels (value is 0 if there is a user file directory only); 3-bit field
EXP_DEV	Device type was explicit
EXP_DIR	Directory specification was explicit
EXP_NAME	File name was explicit
EXP_TYPE	File type was explicit
EXP_VER	Version number was explicit
GRP_MBR	Directory specification is of the group/member number format
HIGHVER	A higher-numbered version (or versions) of the file exists (output from create and enter)
LOWVER	A lower-numbered version (or versions) of the file exists (output from create and enter)
NODE	File specification includes a node name
PPF	File is indirectly accessed process permanent file
QUOTED	File specification includes a quoted string
WILDCARD	File specification string included a wildcard; (this value is returned whenever any of the other wildcard bits are set)
WILD_DIR	Directory specification includes a wildcard
WILD_GRP	Group number contains a wildcard
WILD_MBR	Member number contains a wildcard
WILD_NAME	File name contained a wildcard
WILD_SFD1 through WILD_SFD7	Sub-file directory 1 through 7 specification includes a wildcard
WILD_TYPE	File type contained a wildcard
WILD_UFD	User file directory specification includes a wildcard
WILD_VER	Version number contained a wildcard

CHAPTER 8

RUN-TIME PROCESSING INTERFACE

This chapter describes the interface that VAX-11 RMS uses to access and manipulate files and records within files.

As outlined in Chapter 3, the run-time macro instructions work with the various control blocks to form the record management environment. The file-processing macro instructions deal with the file access block (FAB), and the record-processing macro instructions deal with the record access block (RAB).

The sections that follow discuss the run-time processing interface:

- VAX-11 RMS calling sequence and macro instruction general format
- The path to a file
- Control block usage
- Completion status codes

8.1 THE VAX-11 RMS CALLING SEQUENCE

VAX-11 RMS uses the standard VAX-11 calling sequence and conventions, and preserves all general registers across a call, with the exception of Register0 and Register1. When the routine completes execution, it returns control to the calling program, passing a return status code in Register0. You should analyze the return code to determine the success or failure of the routine and to alter the flow of execution, if necessary.

When you call a VAX-11 RMS routine, you must provide an argument list to define the associated control block (FAB or RAB) and, optionally, any completion routines. The argument list is from two through four longwords in length, as shown in Figure 8-1. (The rename service, however, uses a 5-longword argument list; see Section 13.4.)

RUN-TIME PROCESSING INTERFACE

address for a PUSHAL instruction. If you omit this parameter, no other parameters are allowed; that is, you must use the first format.

The ERR=entry and SUC=entry parameters are optional and, if used, provide the addresses of completion routine entry points. VAX-11 RMS places the values you supply into the argument list on the stack during execution of the expanded macro instruction. These values must be addresses that can be used by a PUSHAL instruction.

When the argument list contains a completion routine argument, the following conventions are used:

- An asynchronous system trap (AST) is queued for the routine when the specified condition (error or success) occurs.
- General registers R0 through R11 are undefined. The argument pointer register (AP) contains the address of the AST argument list (see the VAX/VMS System Services Reference Manual); the AST parameter value in the AST argument list specifies the address of the associated control block (FAB or RAB). The status must be retrieved from the completion status code field (STS) of the associated control block.
- You can modify any general registers saved by an entry mask, in addition to R0 and R1.
- You can issue additional macro instructions for VAX-11 RMS routines within the completion routines.
- To exit from a completion routine, you must perform any necessary clean-up operations and execute a RET instruction.

Note that if the FAB or RAB is invalid, then the error completion routine will not store the error code in the STS field of an invalid structure. The following errors can be detected only by testing Register 0, following the completion of an RMS operation (even if an error completion AST has been specified):

RMS\$_FAB - FAB not writeable or invalid block ID field
RMS\$_RAB - RAB not writeable or invalid block ID field
RMS\$_BLN - invalid block length field (either FAB or RAB)

8.2 THE PATH TO A FILE

Before you can perform operations on a file, you must provide input to the \$OPEN, \$CREATE, \$RENAME, \$PARSE, and \$ERASE macro instructions to establish a path to the file. You do this by setting the file specification string address and size fields (and possibly the default file specification string address and size fields) of the FAB to describe an ASCII string within the program. In this ASCII string, you can have a concatenation of the network node name; a logical or device name; the directory name; and the file name, type, and version number. The following sections describe the processes that resolve all logical names to provide the required file specification components. Appendix C describes the complete file specification syntax and processing algorithms.

RUN-TIME PROCESSING INTERFACE

8.2.1 Interpretation of the File Specification

To establish a path to a file, VAX-11 RMS first calls an internal file specification parse routine. The parse routine forms a fully qualified file specification. If the NAM block specifies an expanded string buffer (see Section 7.2.2), this specification is returned to the user program as the expanded file specification string.

In forming a fully qualified file specification, VAX-11 RMS goes through the following steps:

1. If you specify an open by NAM block (see Section 8.2.3), VAX-11 RMS checks the NAM block fields for a fully qualified file specification.
2. If you do not specify an open by NAM block or if the NAM block fails to provide all the components of a fully qualified file specification, VAX-11 RMS processes the string specified by the file specification string address and string size fields (FNA and FNS) of the FAB. This string may have one of three different forms, which are treated as follows:

- a. If the file specification string has the form

node::"quoted-string"

VAX-11 RMS copies the file specification string without modification to the expanded file specification string.

- b. If the file specification contains only a file name, VAX-11 RMS attempts to translate the string as a logical name. If this attempt succeeds, the equivalence string replaces the original file specification string, and the parse routine restarts. If the attempt fails, the file specification string is taken as the file name and default processing begins (see Section 8.2.2).
- c. If the file specification string has the form
node::device: [directory]file type; version.

If the file name string is neither of the two forms (a and b) discussed above, processing proceeds as follows:

- VAX-11 RMS isolates the various components of the file specification, checks them for correct syntax, and copies them to the expanded file specification string. If the file specification does not include a device name component, default processing begins. If there is a device name component, the VAX-11 RMS parse routine attempts to translate it as a logical name. If a node name has been seen, only user-entered logical names are considered for translation. If the translation attempt fails, the component is treated as a device name.

However, if the translation attempt succeeds, the equivalence string is checked to determine whether it refers to a process permanent file (see the VAX/VMS Command Language User's Guide). If the equivalence string does not refer to a process permanent file, the parse routine restarts, using the equivalence string as its input. If, however, the equivalence string indicates that this is an indirect reference to a process permanent file, the indicated file is therefore the target file resulting from the parse routine, and the logical name is copied to the expanded file specification string.

RUN-TIME PROCESSING INTERFACE

8.2.2 Wild Card Characters in File Specifications

As noted in the VAX/VMS Command Language User's Guide, wild card characters can be used in the last four fields of a file specification. One purpose of wild card characters is to refer to a group of files by a more general file specification, rather than by each of the specific file specifications. There are four characters (or strings of characters) that can be used as wild card characters. These are the asterisk (*), the percent sign (%), the ellipsis (...), and the minus sign (-).

An asterisk is used to match the missing component of a file specification with an alphanumeric character string of any length (including the null string). A percent sign is used to match any single alphanumeric character in that particular position (the null string does not match). The asterisk and the percent sign can be combined in many ways. For example, the sequence:

```
A*E%B*.B*.*
```

matches a file specification or a group of file specifications in which the file name starts with an "A" followed by a string of zero to "n" characters; followed by an "E"; followed by a single character; followed by a "B"; followed by a string of zero to "n" characters. The file type begins with a "B" and is followed by a string of zero to two characters. The version number or numbers in this group will be any and all versions of that file, beginning with the highest version number.

The example file specification matches the following sequence:

```
AEXB.B01  
AZZEYBXXX.B
```

The example file specification does not match the following sequence:

```
AEB.B01  
XAEYB.XBY
```

An asterisk can be used in the following fields of a file specification:

- Directory name
- File name
- File type
- File version number

The percent sign can be used in each of the above fields, with the exception of the file version number field. In this field, only a single asterisk wild card character can be used.

The ellipsis and minus sign wild card characters are aids to searching, or traversing, directory hierarchies. Both the ellipsis and the minus sign allow you to refer to directories in a relative positional sense, rather than by an absolute name for the first directory or group of directories. The ellipsis enables you to select files from all directory levels from a specified level downward to lower levels of the hierarchy. The minus sign, on the other hand, enables you to search up the hierarchy, rather than down. A single minus sign will send the search back up one level from the current default directory level.

RUN-TIME PROCESSING INTERFACE

Wild card characters can be successfully used only in those programs designed to accept them. Most VAX/VMS utilities are designed that way. Prior to using the Open service, it is necessary to use the NAM block and the Parse and Search services to successfully utilize wild card characters as an aid to process groups of files.

8.2.3 File Specification Default Application

If the file specification contains any missing components after VAX-11 RMS completely parses the primary file specification string (specified by the file specification string address and string size fields of the FAB), defaults are applied until either:

- No more components are missing in the specification, or
- No more defaults can be applied.

When VAX-11 RMS applies defaults, program defaults are applied first, in the following order:

1. The default file specification string specified by the contents of the default file specification string address and string size fields (DNA and DNS) of the FAB can supply any of the components necessary to form a full file specification. VAX-11 RMS parses and copies the default file specification string components in the same manner that it does for the primary file specification string (see Section 8.2.1). However, a duplicate field will not cause an error, because VAX-11 RMS ignores any attempt to fill a field that is already occupied.
2. If a NAM block is specified in the FAB, and if a related file NAM block has been specified in the related file field of the NAM block, defaulting can occur as follows. If the related file NAM block has a resultant file specification string, components of the related resultant file specification can be used depending on the state of the OFP (Output File Parse) bit in the file options field (FOP) of the FAB. If the OFP bit is set, VAX-11 RMS parses the output file specification, and only the file name and file type components can be defaulted from the related file (also the file version if the output file version is an explicit wild card character). If the OFP bit is clear (indicating an input file parse), all file specification components, except the file version, are defaulted from the related resultant file specification string.

After program defaults, unless a node specification has been seen, system defaults apply in the following order:

1. If the device name component of the expanded file specification is missing, VAX-11 RMS translates the logical name SYS\$DISK and parses the equivalence string; any expanding components are merged into the expanded file specification string. If the translation yields duplicate components, an error occurs. If the equivalence string includes a logical name, recursion may occur. This step must generate a device name; otherwise, an error occurs.
2. If the directory specification is missing from the expanded file specification, VAX-11 RMS uses the current default directory string from the process I/O control page.

RUN-TIME PROCESSING INTERFACE

After VAX-11 RMS applies the program and system defaults, the expanded name string is complete.

Chapter 13 describes the VAX-11 RMS file specification processing macro instructions. Among the services provided by these macro instructions, the parse service lets you explicitly parse a file specification string independent of the services of the \$OPEN, \$CREATE, \$RENAME, or \$ERASE macro instructions.

8.2.4 Opening and Creating a File by Name Block

When VAX-11 RMS successfully opens a file, the device identification, file identification, and directory identification fields of the NAM block (if present) are filled with the values pertaining to that file.

If you want to reopen the file after it is closed, you can specify the filled-in NAM block by setting the name block address field of the FAB to indicate the address of the NAM block, and setting the NAM option of the file-processing options field (FOP) of the FAB.

If the device identification and file identification fields are nonzero, the file is a fully qualified file specification. However, if the device identification field is nonzero, but the file identification field is 0, VAX-11 RMS uses the normal file specification string parsing routine to supply any missing portions of the full file specification. In this case, the directory specification may come from a nonzero directory identification field of the NAM block. If either the file identification or directory identification field is used, the directory and/or file name, type, and version number of the expanded and resultant file specification strings may be null.

You can create a file the same way that you open a file (above), except that VAX-11 RMS does not use the file identification field as input.

8.3 CONTROL BLOCK USAGE

The control block fields accessed by any run-time macro instruction provide VAX-11 RMS with the means to define or qualify the file and record operations. Depending on the operation, VAX-11 RMS uses one or more of these control blocks with one or more fields being used as input or output to or from the operation. In the chapters that follow, a list of each field being used is provided in the explanation of each macro instruction. Although not individually listed, the block identification (BID) and block length (BLN) fields of every control block used are always inputs to every VAX-11 RMS service.

Before your program calls for the execution of the macro instruction, you must ensure that all the appropriate control block fields used as input contain the necessary values. There are three methods of setting the values in the control block fields:

1. Explicit assembly-time initialization
2. Implicit assembly-time initialization
3. Run-time initialization

RUN-TIME PROCESSING INTERFACE

At assembly time, you explicitly initialize the fields by the use of parameters in the macro instruction for the particular control block (Chapters 4 through 7). You can initialize a field implicitly if VAX-11 RMS has defined a default value for the field. In this case, no action is required on your part. You simply allow the assembly-time expansion of the control block allocation macro instruction to set the default value in the field.

At run time, you can initialize or alter the contents of a control block through the use of the various control block \$xxx STORE macro instructions or directly through instructions that use the defined symbolic offsets associated with the fields (these methods do not provide defaults at run time). If you do not appropriately set a field that is defined as an input field to a particular operation, the operation may fail. VAX-11 RMS assumes that every value found in an input field was placed there for use by the current operation.

8.4 COMPLETION STATUS CODES

Before returning to your program from a file or record operation, VAX-11 RMS indicates the success or failure of the operation by setting a value in the completion status code field (STS) of the associated control block (FAB or RAB).

When first returning to your program after a call to an operation, VAX-11 RMS also sets general register 0 to the value in the status code field. In the case of asynchronous operations, register 0 may simply indicate that the operation is under way.

In the chapters that follow, the discussion of each run-time macro instruction includes a list of the possible nonsevere error and success status codes that you can receive. See Appendix A for a complete list of all VAX-11 RMS status codes.

In general, you may receive one of many error or success codes from an operation. You should thus test for success by checking only the low-order bit of the status code for a true condition (bit is set). The low-order three bits returned in the status code, when taken together, indicate the severity of the code. The severity codes are:

- 001 (1) -- Success (low-order bit set)
- 011 (3) -- Information
- 000 (0) -- Warning; indicates a nonstandard condition. The operation may have performed some, but not all, of the requested function.
- 010 (2) -- Error; you must recognize that a problem exists and provide a contingency plan in your program for such a condition.
- 100 (4) -- Severe error; normally caused by program logic or other unrecoverable condition.

Certain error status codes result in a value's being set in the status value field (STV) of the control block. The description of the codes in Appendix A indicates the instances when the status value field contains such information.

RUN-TIME PROCESSING INTERFACE

When signalling RMS errors, both the STS and the STV fields of the appropriate structure (FAB/RAB) should be supplied. This will cause all relevant information to be displayed in the error message text, including additional information regarding the error status in the STV field. The STV value will be ignored for those errors that do not require it. For example, a GET operation may use an error completion routine to report errors as follows:

```
$GET RAB=MYRAB ERR=REPORT_ERR
:
:
:
.ENTRY REPORT_ERR,0
MOVL 4(AP),R0          ; get structure address into R0
PUSHL RAB$STV (R0)    ; push associated status
PUSHL RAB$STS (R0)     ; push error code
CALLS #2,LIB$SIGNAL    ; signal the error
RET
```

For a more detailed explanation of condition signalling, see the VAX-11 Run-Time Library Reference Manual.

Note that VAX-11 RMS services are considered system services for the purpose of generating system service exceptions on errors (see the VAX/VMS System Services Reference Manual). If you test for error conditions in your program, you should be sure to disable any unwanted system service exception generation.

8.5 PROCESS PERMANENT FILES

A process permanent file is one that is opened (or created) through VAX-11 RMS by supervisor or executive mode code having the PPF bit set in the file processing options (FOP) field of the FAB. This causes the VAX-11 RMS-maintained internal data structures to be allocated in an area of memory in the process control region that remains allocated for the life of the process. Thus, process permanent files can remain open across image activations. You cannot directly access process permanent files by user mode code; you can, however, indirectly access them. VAX-11 RMS provides a subset of the total available operations to the indirect accessor.

Indirect accessors gain access to process permanent files through the logical name mechanism, as follows:

1. The LOGIN command image, or at a later point the command interpreter, opens or creates a file corresponding to the process's input, output, and error message streams. A logical name is created in the process logical name table for SYS\$INPUT, SYS\$OUTPUT, and SYS\$ERROR, respectively. The equivalence string for the logical name has a special format that indicates the correspondence between the logical name and the related process permanent file. For more detail concerning the equivalence-string format for logical names, see the discussion of logical name services in the VAX/VMS System Services Reference Manual. For example, for an interactive user, a single process permanent file is opened for the terminal and all three logical names refer to the one file.

RUN-TIME PROCESSING INTERFACE

2. When an indirect accessor opens or creates a file specifying a logical name that has one of these special equivalence strings, VAX-11 RMS recognizes this and therefore does not open or create a new file; instead, the returned value for the internal file identifier (and later the value for the internal stream identifier from a connect service) is set to indicate that access to the associated process permanent file is with the indirect subset of allowable functions.

Some of the implications for the indirect accessor are:

- A create service for a process permanent file becomes an open service; the fields of the FAB are output according to the description of the open, not the create.
- The open or create service requires no I/O operations.
- Any number of indirect opens and creates are allowed.
- There is only one position context for the file; that is each sequence of the open/create service accesses the same record stream, not an independent stream.
- If the process permanent file was initially opened with the SQO bit set in the file-processing options field, neither random access nor the rewind service is permitted. This is the case for SYS\$INPUT, SYS\$OUTPUT, and SYS\$ERROR.
- Certain options to various services produce errors. For example, you cannot set the NFS, PPF, and UFO bits of the file-processing options field for the open and create services. Other options are ignored, such as the SPL, SCF, and DLT bits of the file processing options field for the close service, the ASY bit of the record-processing options field, and both the multiblock count and multibuffer count fields.
- If a NAM block is used and either an expanded or resultant file specification string is returned, it consists solely of the process logical name followed by a colon, such as SYS\$INPUT:
- The file access field is ignored on an open service; instead, operations are checked against the file access field specified for the original open or create service.
- Information from the record attributes field is saved on each open service (and subsequent connect service) in the value returned in the internal file identifier (and internal stream identifier) field. If the output file is a print file (variable with fixed-control record format and the PRN bit is set in the record attributes field), mapping is performed for each put service from the user-specified carriage control to the print file carriage control format. Thus, different carriage control types from different indirect open services all work correctly.
- You cannot use the erase service.
- Checking is performed for \$DECK, \$EOD, and other dollar sign (\$) records on the SYS\$INPUT stream (see the VAX/VMS Command Language User's Guide).

RUN-TIME PROCESSING INTERFACE

- At image exit time the VAX-11 RMS Rundown control routine insures that the indirect I/O on process permanent files terminates; the process permanent files are not closed.
- You can use only sequential files in this manner.

CHAPTER 9

FILE-PROCESSING MACRO INSTRUCTIONS

VAX-11 RMS provides file-processing macro instructions that you use to perform operations related to the file as a whole. These macro instructions, therefore, deal with fields in the file access block (FAB). See Chapter 4 for a description of the effect of these fields. At run time, the expanded code of these macro instructions causes calls to be made to corresponding VAX-11 RMS services.

In most cases, you use a file-processing macro instruction with parameters to indicate the symbolic address of the FAB and the address of any optional error or success completion routine you may have provided. You can also use the macro instruction without parameters, but you must then create an argument list in your program to define the values for these addresses (see Section 8.1).

Table 3-2 summarizes all the run-time processing macro instructions. This chapter deals only with the following macro instructions, which pertain to file processing:

- \$CLOSE
- \$CREATE
- \$DISPLAY
- \$ERASE
- \$EXTEND
- \$OPEN

For ease of reference, the macro instructions are presented in alphabetical order.

\$CLOSE

9.1 TERMINATING FILE PROCESSING

The \$CLOSE macro instruction invokes the close service, which terminates file processing and closes the file.

You can issue a \$CLOSE macro instruction only when no operation is under way for the file, that is, when all record access blocks (RABs) associated with the file are inactive. Otherwise, the file will not be closed nor will the internal fileidentifier field be set to 0. When the close service operates normally, VAX-11 RMS disconnects all RABs for you, performs the various clean-up procedures (including file

FILE-PROCESSING MACRO INSTRUCTIONS

option and XAB processing), and closes the file. The only types of XABs that the close service processes are the file protection and revision date and time, and then only if the file is write-accessed.

Format

OPERATION	PARAMETERS
label: \$CLOSE	FAB=fab-address ERR=entry SUC=entry

label

Symbolic address for the \$CLOSE macro instruction.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 9-1 lists the FAB fields that VAX-11 RMS uses as input and output for the Close service.

FILE-PROCESSING MACRO INSTRUCTIONS

Table 9-1
Close FAB Fields

Usage	Field Name	Description
Input	FOP	File-processing options (DLT, NAM, RWC, SCF, SPL, and TEF only)
	IFI	Internal file identifier
	NAM	Name block address (used only if NAM is set in file-processing options)
	XAB	Extended attribute block address
Output	IFI	Internal file identifier (zeroed)
	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the close service are listed below. Note that even though a failure may be indicated by the completion status code value, the file was nonetheless closed, if the internal file identifier value was cleared by VAX-11 RMS.

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_ACT ACP file activity precludes operation

RMS\$_DAC File deaccess error

RMS\$_DNR Device not ready

RMS\$_EXP Expiration date not yet reached

RMS\$_MKD File protection violation ACP could not mark file
for deletion

RMS\$_PRV File protection violation

RMS\$_WLK Device write-locked

\$CREATE

9.2 CREATING A FILE

The \$CREATE macro instruction invokes the Create service, which constructs a new file according to the attributes you specify in the FAB. If any extended attribute blocks (XABs) are chained to the FAB, then the characteristics described in the XABs are applied to the file. If an allocation control XAB is present, its allocation quantity (ALQ), allocation options (AOP -- only for the CTG and CBT bits), bucket size (BKZ), and default extension quantity (DEQ) fields are used instead of the corresponding fields of the FAB. When either key definition or allocation XABs are present, they must be densely grouped in ascending order (by REF or AID, respectively). No other types of XABs may intervene. If a name block (NAM) is also connected to the FAB, VAX-11 RMS fills in its fields with information about the created file. The \$CREATE macro instruction leaves the file opened.

The Create service implies PUT access; that is, you need not specify PUT in the file access field of the FAB.

The user should note that the Create-if (CIF) file option (FOP field) specifies simply that: if a file (to be processed) has the same file specification as a file that already exists, then the existing file will be opened and no new file will be created. Some fields in the FAB, such as the file organization (ORG) and record format (RFM) fields, are input to a Create operation, but output from an open operation. For example, the indexed file organization could be specified in the ORG field on a Create-if operation. However, if a sequential file with the same file specification as the indexed file (attempting to be created) already exists, then the existing file will be opened and the ORG field will be set to sequential.

Format

OPERATION	PARAMETERS
label: \$CREATE	FAB=fab-address ERR=entry SUC=entry

FILE-PROCESSING MACRO INSTRUCTIONS

label

A symbolic address for the \$CREATE macro instruction.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 9-2 lists the FAB fields that VAX-11 RMS uses as input and output for the create service.

FILE-PROCESSING MACRO INSTRUCTIONS

Table 9-2
Create FAB Fields

Usage	Field Name	Description
Input	ALQ	Allocation quantity. This field is ignored if an allocation XAB is present.
	BKS	Bucket size. This field is ignored if an allocation XAB is present.
	BLS	Block size. (magnetic tape only)
	DEQ	Default file extension quantity. This field is ignored if an allocation XAB is present.
	DNA	Default file specification string address
	DNS	Default file specification string size
	FAC	File access
	FNA	File specification string address
	FNS	File specification string size
	FOP	File-processing options (DLT, NAM, RWC, SCF, SPL, and TEF only)
	FSZ	Fixed control area size
	IFI	Internal file identifier (must be 0)
	MRN	Maximum record number (relative organization only)
	MRS	Maximum record size
	NAM	Name block address
	ORG	File organization
	RAT	Record attributes
	RFM	Record format
	RTV	Retrieval window size
	SHR	File sharing
XAB	Extended attribute block address	
Output	ALQ	Allocation quantity (contains actual number of blocks allocated)
	BLS	Block size (sequential organization only)
	DEV	Device characteristics
	IFI	Internal file identifier
	SDC	Spooling device characteristics
	STS	Completion status code (also returned in Register 0)
	STV	Status value (contains the I/O channel number if the operation is successful)

FILE-PROCESSING MACRO INSTRUCTIONS

Table 9-3 lists the NAM block fields that VAX-11 RMS uses as input and output for the create service if the name block address field is specified in the FAB.

Table 9-3
Create NAM Block Fields

Usage	Field Name	Description
Input	DID	Directory identification (input only if NAM bit is set in the file processing options (FOP) field of FAB)
	DVI	Device identification (input only if NAM bit is set in the FOP field of the FAB)
	ESA	Expanded string area address
	ESS	Expanded string area size
	RLF	Related file NAM block address (if nonzero, RSA and RSL are input from related file NAM block)
	RSA	Resultant string area address
	RSS	Resultant string area size
Output	DID	Directory identification
	DVI	Device identification
	ESL	Expanded string length (if, on input, both the ESA and ESS are nonzero, and if the NAM bit of the FOP field of the FAB is clear or DID is 0, the expanded file specification string is copied to the buffer specified by the input ESA field)
	FID	File identification
	FNB	File name status bits (FNB is output only if NAM bit in FOP field of FAB is clear, or if DID field was 0 on input)
	RSL	Resultant string length (if RSA and RSS are both nonzero on input, the resultant file specification is copied to the buffer specified by RSA)

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Create service are listed below. If a failure is indicated, the file may indeed have been created, but will not be opened for processing, depending on the nature of the failure.

FILE-PROCESSING MACRO INSTRUCTIONS

Success:

RMS\$_CREATED File was created, not opened. This status is returned when the CIF option is used and the file must be created. If the file is opened, RMS\$_NORMAL is returned.

RMS\$_NORMAL Operation successful

RMS\$_SUPERSEDE Created file supersedes an existing file

Failure:

RMS\$_ACT File activity precludes operation

RMS\$_CRE ACP file create error

RMS\$_DNF Directory not found

RMS\$_DNR Device not ready

RMS\$_EXP Expiration date not yet reached

RMS\$_FEX File already exists

RMS\$_FLK File locked; not available

RMS\$_PRV File protection violation

RMS\$_WLK Device write-locked

\$DISPLAY

9.3 OBTAINING ATTRIBUTES OF A FILE

The \$DISPLAY macro instruction invokes the Display service, which retrieves file attribute information about a file and places this information in fields in the XABs chained to the FAB. VAX-11 RMS determines the type of file attribute information needed by the type of XABs present. Prior to invoking the Display service, the file must already have been opened for access by a Create or Open operation.

Format

OPERATION	PARAMETERS
label: \$DISPLAY	FAB=fab-address ERR=entry SUC=entry

FILE-PROCESSING MACRO INSTRUCTIONS

label

A symbolic address for the \$DISPLAY macro instruction; optional.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 9-4 lists the FAB fields that VAX-11 RMS uses as input and output for the Display service.

Table 9-4
Display FAB Fields

Usage	Field Name	Description
Input	IFI	Internal file identifier
	XAB	Extended attribute block address
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value; contains the address of the XAB that caused error.

VAX-11 RMS places the attribute values in the corresponding fields of the appropriate XAB.

Note that the Open service performs an implicit Display service (see Section 9.6).

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the display service are listed below.

Success:

RMS\$_NORMAL Operation successful

RMS\$_OK_NOP XAB not filled in when file opened for block I/O.

FILE-PROCESSING MACRO INSTRUCTIONS

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_PRV	File protection violation

\$ERASE

9.4 DELETING A FILE

The \$ERASE macro instruction invokes the Erase service, which deletes a VAX-11 RMS disk file and removes the file's directory entry as specified in the path to the file (see Section 8.2). You must use the \$REMOVE macro instruction to delete additional directory entries, if any (see Chapter 13).

Deleting a file releases the file's allocated space for use by another file; the deletion does not physically remove the data (as does overwriting or zeroing). Only files that are closed can be deleted; an open file cannot be deleted with the erase service, but may be deleted by the \$CLOSE macro instruction by setting the DLT bit in the file-processing options field of the FAB. Furthermore, you cannot delete files from magnetic tape volumes.

Format

OPERATION	PARAMETERS
label: \$ERASE	FAB=fab-address ERR=entry SUC=entry

label

A symbolic address for the \$ERASE macro instruction; optional.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

FILE-PROCESSING MACRO INSTRUCTIONS

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 9-5 lists the FAB fields that VAX-11 RMS uses as input and output for the Erase service.

Table 9-5
Erase FAB Fields

Usage	Field Name	Description
Input	DNA	Default file specification string address
	DNS	Default file specification string size
	FNA	File specification string address
	FNS	File specification string size
	FOP	File-processing options (NAM bit only)
	IFI	Internal file identifier (must be 0)
	NAM	Name block address
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value

Table 9-6 lists the NAM block fields that VAX-11 RMS uses as input and output for the erase service if the name block address field is specified in the FAB.

FILE-PROCESSING MACRO INSTRUCTIONS

Table 9-6
Erase NAM Block Fields

Usage	Field Name	Description
Input	DID	Directory identification (input only if NAM bit is set in the file processing options (FOP) field of FAB)
	DVI	Device identification (input only if NAM bit is set in the FOP field of the FAB)
	ESA	Expanded string area address
	ESS	Expanded string area size
	FID	File identification (input only if NAM bit is set in the FOP field of the FAB)
	RLF	Related file NAM block address (if nonzero, RSA and RSL are from related file NAM block)
	RSA	Resultant string area address
	RSS	Resultant string area size
Output	DID	Directory identification
	DVI	Device identification
	ESL	Expanded string length (if, on input, both the ESA and ESS are nonzero, and if the NAM bit of the FOP field of the FAB is clear or DID is 0, the expanded file specification string is copied to the buffer specified by the input ESA field)
	FNB	File name status bits
	RSL	Resultant string length (if RSA and RSS are both nonzero on input, the resultant file specification is copied to the buffer specified by RSA)

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Erase service are listed below.

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_DNF Directory not found

RMS\$_DNR Device not ready

RMS\$_MKD ACP could not mark file for deletion

RMS\$_PRV File protection violation

RMS\$_WLK Device write-locked

\$EXTEND

9.5 EXTENDING A FILE'S ALLOCATED SPACE

The \$EXTEND macro instruction invokes the Extend service, which increases the amount of space allocated to a VAX-11 RMS disk file. You can only extend open files; otherwise, an error occurs.

The allocation quantity field of the FAB (or the allocation XAB, if used) must contain the number of blocks that VAX-11 RMS is to add to the file. Furthermore, you can indicate other attributes regarding the manner and location for allocation. For example, you can indicate that the additional blocks must be allocated contiguously. If you do, however, and not enough contiguous space is available, the operation will fail. (This extension does not have to occur contiguous to the initial file space.)

If an allocation control XAB is present, its allocation quantity (ALQ) and allocation options (AOP -- the CBT and CTG bits only) fields are used instead of the corresponding fields in the FAB. The allocation quantity field of the XAB is set to the actual extension size.

Format

OPERATION	PARAMETERS
label: \$EXTEND	FAB=fab-address ERR=entry SUC=entry

label

A user-defined symbolic address for the \$EXTEND macro instruction; optional.

Fab=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 9-7 lists the FAB fields that VAX-11 RMS uses as input and output for the Extend service.

FILE-PROCESSING MACRO INSTRUCTIONS

Table 9-7
Extend FAB Fields

Usage	Field Name	Description
Input	ALQ	Allocation quantity. This field is ignored if an allocation XAB is present.
	FOP	File-processing options. Checked to see if the CTG or CBT bit is set to indicate contiguous allocation; ignored if allocation XAB is present.
	IFI	Internal file identifier
	XAB	Extended attribute block address. Only the allocation type of XAB will be processed.
Output	ALQ	Allocation quantity (contains the actual extension allocation value if no allocation XAB is present)
	STS	Completion status code (also returned in Register 0)
	STV	Status value (contains the total of blocks allocated, totaled across all allocation XABs)

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the extend service are listed below.

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_ACT File activity precludes operation

RMS\$_DNR Device not ready

RMS\$_EXT ACP file extend error

RMS\$_WLK Device write-locked

\$OPEN

9.6 OPENING AN EXISTING FILE

The \$OPEN macro instruction invokes the open service, which makes an existing file available for processing by your program. This macro instruction implements the type of access desired, and sets the degree to which the file can be shared. You must open a file before you perform any record operations. If any XABs are chained to the FAB, VAX-11 RMS places the attribute values in the fields of the appropriate XAB. If you specify a NAM block in the FAB, the contents of the device, directory, and file identification fields can be used

FILE-PROCESSING MACRO INSTRUCTIONS

to perform an open by NAM block (see Section 8.2.3). In addition, the various fields of this NAM block are filled in with auxiliary file specification information.

Format

OPERATION	PARAMETERS
label: \$OPEN	FAB=fab-address ERR=entry SUC=entry

label

A symbolic address for the \$OPEN macro instruction; optional.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 9-8 lists the FAB fields used as input and output for the open service.

FILE-PROCESSING MACRO INSTRUCTIONS

Table 9-8
Open FAB Fields

Usage	Field Name	Description
Input	DEQ	Default file extension quantity. If a nonzero value is present in this field, it applies to this open of the file only.
	DNA	Default file specification string address
	DNS	Default file specification string size
	FAC	File access
	FNA	File specification string address
	FNS	File specification string size
	FOP	File-processing options (see Section 4.2.14)
	FSZ	Fixed control area size; unit record devices only.
	IFI	Internal file identifier (must be 0)
	NAM	Name block address
	RAT	Record attributes; unit record devices only
	RFM	Record format; unit record devices only
	RTV	Retrieval window size
	SHR	File sharing
XAB	Extended attribute block address	
Output	ALQ	Allocation quantity; contains the highest numbered block allocated to the file.
	BKS	Bucket size; not used for sequential files
	BLS	Block size; for sequential files only
	DEQ	Default file extension quantity
	DEV	Device characteristics
	FOP	File-processing options; the bits CTG, CBT, RCK, and WCK are set or cleared individually according to the file attributes
	FSZ	Fixed control area size; only applies to variable with fixed length control records
	IFI	Internal file identifier
	MRN	Maximum record number; for relative files only
	MRS	Maximum record size
	ORG	File organization
	RAT	Record attributes
	RFM	Record format
	SDC	Spooling device characteristics
	STS	Completion status code (also returned in Register 0)
STV	Status value (contains the I/O channel number if the operation is successful)	

FILE-PROCESSING MACRO INSTRUCTIONS

Table 9-9 lists the NAM block fields (see Chapter 7) that VAX-11 RMS uses as input and output for the Open service, if the name block address field is specified in the FAB.

Table 9-9
Open NAM Block Fields

Usage	Field Name	Description
Input	DID	Directory identification (input only if NAM bit is set in the file processing options (FOP) field of FAB)
	DVI	Device identification (input only if NAM bit is set in the FOP field of the FAB)
	ESA	Expanded string area address
	ESS	Expanded string area size
	FID	File identification (input only if NAM bit set in FOP field of FAB)
	RLF	Related file NAM block address (if non-zero, RSA and RSL are from related file NAM block)
	RSA	Resultant string area address
	RSS	Resultant string area size
Output	DID	Directory identification
	DVI	Device identification
	ESL	Expanded string length (if, on input, both the ESA and ESS are nonzero, and if NAM bit of the FOP field of the FAB is clear or DID and FID are 0, the expanded file specification string is copied to the buffer specified by ESA)
	FID	File identification
	FNB	File name status bits
	RSL	Resultant string length (if RSA and RSS are both nonzero and if NAM bit is clear or FID is 0, the resultant file specification is copied to the buffer specified by RSA)

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the open service are listed below.

Success:

RMS\$_KFF	Known file found
RMS\$_NORMAL	Operation successful
RMS\$_OK_NOP	XAB not filled in when file opened for block I/O

FILE-PROCESSING MACRO INSTRUCTIONS

Failure:

RMS\$_ACC	ACP file access error
RMS\$_ACT	File activity precludes operation
RMS\$_DNF	Directory not found
RMS\$_DNR	Device not ready
RMS\$_FLK	File locked; not available
RMS\$_FNF	No such file exists
RMS\$_PRV	File protection violation
RMS\$_WLK	Device write-locked

CHAPTER 10

RECORD OPERATION PERFORMANCE

Some of the key concepts that you must understand in relation to record operations are:

- Record access
- Current record context
- Record streams
- Synchronous and asynchronous operations
- Record locking

Sections 10.1 through 10.5 discuss these concepts; Chapter 11 describes each record-processing macro instruction in detail.

10.1 RECORD ACCESS

To process a record, you must identify the record and specify the record access mode you are going to use. Once the record is identified, you have two different record transfer modes available to manipulate it. The following sections describe how you specify the record access mode and the transfer mode.

10.1.1 Specifying the Record Access Mode

The value that you set in the record access mode field of the RAB tells VAX-11 RMS what type of record access to use for the particular record operation. During program execution, you can switch the record access mode by changing the contents of this field. This is known as dynamic access.

VAX-11 RMS lets you set any one of the following three values:

1. SEQ -- this value indicates the sequential record access mode. When you use this record access mode, the access will be a function of the Next Record (see Section 10.2), and no additional record specification is necessary. This record access mode is valid for any file organization.

RECORD OPERATION PERFORMANCE

2. KEY -- this value indicates random access by key. This record access mode is used with relative files and sequential files on disk with fixed-length records to denote random access by relative record number and with indexed files to denote random access by key value. The key value for the record to be found or retrieved is placed in the key buffer, which is described by the values set in the key buffer address and key size fields of the RAB. When accessing an indexed file, the particular key of reference (index to search on) must be specified in the KRF field of the RAB.
3. RFA -- this value indicates that access is random by the record's file address (RFA). This record access mode is limited to retrieval operations for disk files.

To use this access mode, you must save the RFA that VAX-11 RMS returned from a previous operation. Then, before you initiate a new operation, you specify access by RFA mode in the record access mode field of the RAB, and restore the RFA. The RFA does not change when you close a file and later reopen it.

The format of the RFA is known internally to VAX-11 RMS.

VAX-11 RMS examines the contents of the record access mode field of the RAB during the execution of a \$GET, \$FIND, or \$PUT macro instruction. You need not specify a record access mode for operations with a \$UPDATE, \$DELETE, or \$TRUNCATE macro instruction. However, you cannot request these operations until you have first accessed the target record with a \$GET or \$FIND macro instruction.

10.1.2 Specifying the Record Transfer Mode

The record-processing option field of the RAB lets you specify the record transfer mode. There are two record transfer modes -- locate and move -- which tell VAX-11 RMS how to access the target record for the get service (\$GET macro instruction) once the record is in memory. You can switch the record transfer mode while your program is executing by changing the contents of the record-processing option field.

In the record-processing option field you indicate locate mode by setting the LOC bit. If you do not set this bit, VAX-11 RMS uses move mode, by default.

In locate mode, your program accesses records directly in an I/O buffer. Therefore, VAX-11 RMS normally does not need to move records between I/O buffers and a user program buffer. VAX-11 RMS does not support locate mode for operations involving the \$PUT or \$UPDATE macro instructions. However, the \$GET macro instruction supports locate mode operations on files of all organizations. Note that locate mode, even if specified, may not actually be used due to the occurrence of any of the following:

1. Records crossing block boundaries
2. The file access field of the FAB being set to UPD
3. Multiple record streams
4. Indirect accessor of process - permanent files (see Section 8.5)

RECORD OPERATION PERFORMANCE

In move mode, VAX-11 RMS transfers individual records between I/O buffers and your program buffer. For the \$GET macro instruction, VAX-11 RMS reads a block (for sequential files) or a bucket (for relative and indexed files) into an I/O buffer. VAX-11 RMS then selects the desired record from the buffer and moves it to a program-specified location.

When writing records to the file (\$PUT and \$UPDATE), your program first builds a record in any desired program location, stores its address and size in the RAB, and calls the appropriate VAX-11 RMS routine as specified by the particular macro instruction. VAX-11 RMS moves the record from its specified location into an I/O buffer. Depending upon the file organization and options, the buffer may be written immediately or only when it is filled.

10.2 CURRENT RECORD CONTEXT

For each RAB connected to a file access block (FAB), VAX-11 RMS maintains current context information, identifying where each RAB is positioned at any point in time. VAX-11 RMS modifies the current context as your program performs record operations.

At any point in time, the current context is represented by, at most, two records:

1. The Current Record
2. The Next Record

The context of these two records is internal to VAX-11 RMS; you have no direct contact with them. However, an explanation of their purpose and importance can aid in your understanding of how VAX-11 RMS works.

10.2.1 Current Record

The Current Record represents the target record for \$UPDATE, \$DELETE, or \$TRUNCATE macro instructions. The Current Record also facilitates sequential processing on disk devices for a stream. VAX-11 RMS rejects any update, delete, or truncate request, if there is no current record defined. In addition, an operation with a \$GET macro instruction using sequential record access mode and immediately preceded by a \$FIND macro instruction operates on the record specified by the Current Record. If the find service did not lock the record (for relative and indexed file organizations) and the current record has been deleted, the get service will access the next existing record.

When a RAB is first connected to a FAB, the Current Record is undefined. Furthermore, any unsuccessful record operation, or successful execution of a macro instruction other than \$GET or \$FIND, causes the Current Record to be undefined.

RECORD OPERATION PERFORMANCE

The Current Record is set to the RFA of the record upon which an operation is performed with a \$GET or \$FIND macro instruction. VAX-11 RMS also places this address in the record's file address field of the RAB. This means that:

1. After initialization, the Current Record always refers to the record's file address of the most recent successful operation with a \$GET or \$FIND macro instruction (unless failure occurs or a macro instruction other than \$GET or \$FIND executes).
2. The record's file address field of the RAB, unless you modify it, always contains the address of the target record (if the operation fails, the record's file address is undefined).

Table 10-1 summarizes the effect that each successful record operation has on the context of the Current Record.

10.2.2 Next Record

VAX-11 RMS uses the Next Record for operations involving sequential record access mode. When the record access mode field of the RAB indicates sequential processing, the Next Record represents the target record for the next operation involving:

- The \$FIND macro instruction
- The \$PUT macro instruction
- The \$GET macro instruction (if the immediately preceding operation was not a \$FIND macro instruction); if the next record cell in a relative file organization does not contain a record, the target record is the next existing record.

This "look-ahead" ability significantly decreases access time for sequential processing. VAX-11 RMS uses its internal knowledge of file organization and structures to determine the Next Record as follows:

- Operations with the \$CONNECT macro instruction initialize the Next Record to:
 - The first record or cell in a file of sequential or relative organization, respectively
 - The first record in the collating sequence of the specified key of reference in an indexed file.
 - The end of a sequential file on disk if the record processing options field of the RAB has the EOF option bit set.
 - The end of a write-accessed magnetic tape file unless the file processing options field of the FAB has the NEF bit set.
- Operations with the \$GET macro instruction in any record access mode and the \$FIND macro instruction in sequential record access mode cause the Next Record to indicate the next record or cell in the file.

RECORD OPERATION PERFORMANCE

- Operations with a \$TRUNCATE macro instruction cause the Next Record to indicate the end of file. Therefore, you need only use \$PUT macro instructions after truncation to extend the file. You can truncate only sequential files.
- Operations with the \$FIND or \$PUT macro instructions in random access mode have no effect on the Next Record.
- Operations with the \$PUT macro instruction in sequential access mode initialize the Next Record to:
 - The end of file in a sequential file.
 - The next record or cell in a relative file.
- Operations with the \$PUT macro instruction in sequential access mode in an indexed file cause the Next Record to be undefined.
- Operations with the \$DELETE, \$UPDATE, \$FREE, or \$RELEASE macro instructions in any record access mode have no effect on the Next Record.
- Operations with the \$REWIND macro instruction in any record access mode cause the Next Record to indicate the first record or cell in the file.
- Any unsuccessful record operation has no effect on the Next Record.

Table 10-1 summarizes the effect that each successful record operation has on the Next Record.

10.3 RECORD STREAMS

Before you can process the records in a file, you must first establish a record stream to that file. A record stream is the logical association of a RAB with a FAB. Once you have established this association, you can issue requests for operations on the records in the file that the FAB represents.

For all but the sequential file organization, there can be any number of RABs associated with a single FAB, and each RAB represents an independent record stream. Sequential files with 512-byte fixed-length records can also have multiple streams. If you establish a single record stream, your program uses the stream to issue a sequence of record operations, which are executed serially. Therefore, you can process only one record at a time. However, when you establish multiple record streams for a file, you can process a record from each stream in parallel. Therefore, multiple record streams provide concurrently active sequences of record operations to the same file.

After you open a file by issuing a \$OPEN (or \$CREATE) macro instruction, you establish the record stream by placing the address of the FAB in the file access block field of the appropriate RAB or RABs. Then, you issue a \$CONNECT macro instruction. Once you have completed the desired sequence of operations, you terminate the association by issuing a \$DISCONNECT macro instruction.

Chapter 11 describes the \$CONNECT and \$DISCONNECT macro instructions.

RECORD OPERATION PERFORMANCE

Table 10-1
Record Access Stream Context

Record Operation	Record Access Mode	Current Record	Next Record
Connect	does not apply	none	first record
Connect with EOF bit set in record-processing options field	does not apply	none	end of file (Does not apply to indexed files)
Get last operation not a find	sequential	new	new Current Record+1
Get last operation was a find	sequential	unchanged	Current Record+1
Get	random	new	new Current Record+1
Put	sequential	none	1. sequential file-- end of file 2. relative file-- next record position 3. indexed file-- undefined
Put	random	none	unchanged
Find	sequential	new	new Current Record+1
Find	random	new	unchanged
Update	does not apply	none	unchanged
Delete	does not apply	none	unchanged
Truncate	does not apply	none	end of file
Rewind	does not apply	none	first record
Free	does not apply	none	unchanged
Release	does not apply	none	unchanged

NOTES:

1. Except for the truncate operation, VAX-11 RMS establishes the Current Record before establishing the identity of the Next Record.
2. The notation "+1" indicates the next sequential record as determined by the file organization. For indexed files, the current key of reference is part of this determination.
3. The connect operation on an indexed file establishes the Next Record to be the first record in the index represented by the RAB key of reference (KRF) field.
4. The connect operation leaves the Next Record as the end of file for a magnetic tape file opened for put operations (unless the NEF bit is set).

RECORD OPERATION PERFORMANCE

10.4 SYNCHRONOUS AND ASYNCHRONOUS OPERATIONS

Within each record stream, VAX-11 RMS lets you perform operations either synchronously or asynchronously. In synchronous operations, VAX-11 RMS returns control to your program only after the record operation request is satisfied.

In asynchronous operations, VAX-11 RMS may return control to your program before the operation is satisfied. In this way, your program can use the time required to transfer data between the file and memory to perform other computations. Note that in asynchronous operations, the operation may complete before control is returned to your program. This is due to several factors. For example, the required record may already reside in an I/O buffer, or the operating system may schedule another program, thus possibly allowing a necessary operation to complete before the original program is rescheduled.

Generally, VAX-11 RMS executes in either executive mode or executive AST mode. When a VAX-11 RMS operation is initiated, processing begins in executive mode. If device I/O is necessary to process the request, the QIO system service is called. VAX-11 RMS specifies an AST to signal completion. At this point, VAX-11 RMS will exit from executive mode. If the operation is being performed asynchronously, control is returned to the caller. If the operation is synchronous, VAX-11 RMS waits for an event flag in the access mode of the caller. When the I/O is complete, VAX-11 RMS will continue processing in executive AST mode. Thus user-mode ASTs can be serviced while a synchronous VAX-11 RMS operation called from user mode is awaiting I/O completion. However, processing in user mode during an asynchronous VAX-11 RMS operation will be interrupted by VAX-11 RMS processing in executive AST mode when I/O completes.

VAX-11 RMS should not be called from kernel mode. Nor should it be called from executive mode when executive-mode ASTs are disabled.

The following sections describe how to declare synchronous and asynchronous operations.

10.4.1 Synchronous Operations

To declare a synchronous operation, you must clear the ASY bit in the record-processing options field of the RAB. Since by default this bit is off at assembly time, you normally do not have to set it off unless you had set it on previously.

Normally, you would not use success and error routines with synchronous operations. Instead, you would test the completion status code for an error and change the program's flow accordingly. However, if you use these routines, they will be executed as asynchronous system traps (ASTs) before the inline return to your program (unless ASTs are disabled).

As explained in Section 10.4, user-mode AST routines may be executed before the completion of a synchronous record operation. If an AST routine attempts to perform operations on a file which is being called from a non-AST level, it must be prepared to handle stream-activity errors (RMS\$_RSA) as discussed in the next section.

RECORD OPERATION PERFORMANCE

10.4.2 Asynchronous Operations

To declare an asynchronous record operation, you must set the ASY bit in the record-processing options field of the RAB. You can switch between synchronous and asynchronous operations during processing of a record stream by setting or clearing the ASY bit on a per-operation basis.

You can specify completion routines to be executed as ASTs if success or error conditions occur. Within such routines, you can issue additional operations, but they too should be asynchronous. Otherwise, all other currently active asynchronous requests in your program cannot have their completion routines executed until the synchronous operation completes.

If an asynchronous operation is not yet complete at the time of return from a call to a VAX-11 RMS service, the completion status field of the RAB will be 0, and a success status code of RMS\$ PENDING will be returned in Register 0. This status code indicates that the operation was initiated but is not yet complete. You must never modify the contents of a RAB when an operation is in progress.

If you issue a second record operation request for the same stream before a prior request is complete, you will receive an error status code of RMS\$ RSA, indicating that the record stream is still active. This can also occur when an AST level routine attempts to use an active record stream; the original I/O request may be synchronous or asynchronous. In either case it is the caller's responsibility to recognize the possibility and prevent the problem by issuing a \$WAIT macro instruction (see Chapter 11).

Note that the connect operation may be performed asynchronously. If the ASY option is set at assembly time, a \$WAIT macro instruction should follow the \$CONNECT in order to synchronize with the completion of the operation. Another technique would be to perform the connect operation synchronously and set the ASY option only at run time, after the connect operation.

Upon completion of the operation, your program receives control at the point following the \$WAIT macro instruction.

10.5 FILE SHARING

VAX-11 RMS file sharing allows multiple access streams to a single file which enables the user to concurrently read, write and modify records within the file, in a controlled manner. Independent record access streams are associated with a single file in one or both of the following ways:

1. Multiple record access blocks (RABs) are associated with a single file access block (FAB), using the connect operation. This is known as multi-streaming.
2. Multiple file access blocks (FABs) specify that the same file will be associated with that file when opened. One or more record access blocks (RAB's) may then be connected.

The sharing field (SHR) in the FAB is used to declare whether multiple streams and/or multiple openers of the same file will be allowed. Enabling the multi-stream option (MSE) in the SHR field will allow more than one RAB to be connected to the FAB.

RECORD OPERATION PERFORMANCE

When multiple streams are connected, the buffers allocated for each stream become part of a buffer cache for the entire file. A record operation on one stream may use cached buffers from a previous record operation that referenced the same buckets.

In either form of sharing, VAX-11 RMS controls the reading and writing of I/O buffers to ensure file and record structure integrity. All relative and indexed files, and sequential files with 512-byte fixed-length records, can be shared in this manner. The RMSSHARE utility is used to initialize this shared file database. This utility is run as part of the system start-up procedure. If the size of the shared file database is inadequate for your system, sharing-pagecount exceeded (RMS\$SPE) or dynamic memory errors (RMS\$DME) may be encountered while processing shared files. See the VAX/VMS Systems Manager's Guide for documentation on using RMSSHARE to monitor and modify the size of the shared file database.

All sequential files may be write-shared with user provided interlocks. To use this feature, you must set the UPI bit in the SHR field of the FAB. When the UPI option is used, VAX-11 RMS does not attempt to control the reading and writing of I/O buffers across processes; nor does it maintain end-of-file information. The \$FLUSH macro instruction is used to force the writing of modified I/O buffers and is also used to rewrite the record attributes (including end of file) in the file headers. Processes that open the file after that point will obtain the new end-of-file information. Note also that record attributes are rewritten whenever a file is closed. The last write accessor to close the file must also be the last accessor to have extended the file. If not, end-of-file information will be written by another write accessor. Read accessors of a shared sequential file can update their internal end-of-file context by closing and reopening the file.

The UPI form of sharing is applicable only to sequential files. In all other cases, VAX-11 RMS transparently controls the reading and writing of buffers to the file, and always maintains current end-of-file information. The user need not be concerned with all of the issues mentioned above. The record locking facility, described in the following sections, controls access to individual records within the file. Automatic record locking is enabled whenever the file is shared (except for UPI sharing on sequential files) and requires only that the user be prepared to handle record-lock errors (RMS\$_RLK) on Find or Get operations.

10.6 RECORD LOCKING

VAX-11 RMS provides a record-locking capability for relative files, indexed files, and sequential files with 512-byte fixed-length records. This capability affords control over operations when more than one stream or process is simultaneously accessing the file. Record locking makes certain that when a program is adding, deleting, or modifying a record on a given stream, another stream or process cannot access the same record.

Record locking occurs on a file accessed for some form of writing (FAC is set to either PUT, UPD, or DEL) only if the file-sharing field (SHR) of the FAB is set to some form of writing or the MSE bit is set.

There are two types of record locking: automatic and manual. VAX-11 RMS handles automatic record locking transparently. You use it when you are dealing with a lock on a single record at a time. Manual record locking requires additional effort on your part. You use it when dealing with locks on multiple records at one time.

RECORD OPERATION PERFORMANCE

A record can be in any of three states: unlocked, automatically locked, or manually locked. When a record is initially locked, it is in either the manually or automatically locked state. It will remain in that state until the lock is released. That is, it cannot move directly from the automatically to manually locked state, or vice versa. Therefore, you make an initial decision based on your needs to use automatic or manual locking for a given record, and continue to use the same type of locking with that record until the lock is released.

The following sections describe the two types of record locking.

10.6.1 Automatic Record Locking

For automatic record locking, the lock occurs on every execution of a \$FIND or \$GET macro instruction (unless the NLK bit is set in the record-processing options field). The lock is held until the next operation on the stream; that is, the lock is released when the Next Record is accessed, the Current Record is updated or deleted, the record stream is disconnected, the file is closed, or an operation causing an error occurs. Therefore, the record is freed when you issue any of the following macro instructions:

- \$FIND
- \$GET
- \$PUT
- \$UPDATE
- \$DELETE
- \$REWIND
- \$DISCONNECT
- \$CLOSE
- \$FREE
- \$RELEASE

The \$FREE and \$RELEASE macro instructions let you explicitly unlock the record.

If you place a record in an empty cell in a relative file with a \$PUT macro instruction, the cell is, in effect, locked by the put service. It is unlocked when the service completes.

One exception to the automatic unlocking exists: a record remains automatically locked on a sequential GET service following a find service that caused the record to be locked. The automatic record-locking scheme normally does exactly what is required without any interruption. For example, to update an existing record, the following sequence could be used:

```
$GET RAB=A_RAB
.
. (code to modify the record buffer)
.
$UPDATE RAB=A_RAB
```

RECORD OPERATION PERFORMANCE

The \$GET macro instruction reads the record into a buffer in order to examine and modify the record. It also establishes the current record and locks it in preparation for the \$UPDATE macro instruction. The program then operates on the record as required. When the record is finally updated, the record lock is released. During the time that elapses between the Get and Update operations, other streams attempting to access that same record will receive a record-lock error (RMS\$RLK). This will prevent the original record from being accessed and potentially modified, before the stream has finished operating on it. When the record lock is released by the update operation, the modified record will be accessible by other streams.

10.6.2 Manual Record Locking

For manual record locking, you have explicit control over the unlocking of records. Thus, manual record locking lets you control operations that must be done together.

Manual record locking occurs when the ULK bit is set in the record-processing options field on the execution of a \$GET, \$FIND, or \$PUT macro instruction. (These three macro instructions will also unlock any record that was locked with automatic record locking.) Once the record is manually locked, it will remain in that state until explicitly unlocked by either the free or release service, or until the stream terminates (by a disconnect or close service). Other operations on the record or stream, including operations that result in errors, do not cause the record to be unlocked.

Manual control over the unlocking of records is useful when multiple records must be modified as a single transaction. An example of this would be a case in which two separate records are randomly accessed and updated. The first record must not be accessed by another stream until modifications to the second record are complete. The program attempts to update the first record, but at the same time retains the lock on it. Thus, in the event of a failure to update the second record, the original contents of the first record could be restored. Manual unlocking is specified when accessing the first record. This will prevent the record from being automatically unlocked after the update operation. The lock is released by using the \$FREE macro instruction after successively updating the second record (the normal case), or after restoring the original contents of the first record (the error condition). The \$FREE macro instruction releases all locks for that stream, simultaneously. At this time, the updated records will become accessible to other streams.

The \$RELEASE macro instruction is used to selectively release manually locked records. It unlocks the record, using its record file address (RFA).

RECORD OPERATION PERFORMANCE

10.6.3 Controlling Record Locking

Three of the bits in the record-processing options field (ROP) of the RAB control manual record locking and unlocking. As in the case of the ULK bit, the following are input to \$FIND, \$GET, or \$PUT (with the ULK) macro instructions. They are:

1. NLK - do not lock record
2. RLK - lock record, but allow readers
3. NXR - lock non-existent record (not applicable to indexed sequential files)

The NLK bit specifies that the record accessed with either a \$GET or \$FIND macro instruction is not to be locked. Specifying only get access (FAC field) for the file also implies that records are not to be locked. In either case, if the target record is locked by another stream, a record-locked error (RMS\$_RLK) will be returned. The only exception to this will be if the stream locking the record has allowed readers (see RLK below). Records accessed for purposes other than modifying (i.e., deleting or updating) should not be locked. This will reduce the probability of other streams receiving record-lock errors. Attempting to delete or update a record that was not locked will fail. The NLK bit takes precedence over the ULK bit described above.

Streams that are locking records for modification may allow non-locking streams (as described above) to read locked records. The RLK bit specifies that the record will be locked for possible modifications. However, readers will be able to access it. When a non-locking stream reads a record locked in this manner, an alternate success code indicates that the target record was indeed locked, but that readers are allowed by the locker (RMS\$_OK_RLK). Another stream attempting to lock the record, however, will still receive a record-locked error (RMS\$_RLK).

The non-existent record lock (NXR) option applies only to relative files, or to sequential files with 512-byte fixed-length records. It is used to lock randomly accessed records that do not already exist in the file at the time of access. This will prevent other streams from putting a new record into that cell until the stream that locked it either: puts a record there itself, or releases the record lock. For example, suppose that a file contains records one through ten. A program attempting to randomly access record fifteen would normally receive either a record-not-found (RMS\$_RNF) error if the file was a relative file, or an end-of-file (RMS\$_EOF) error if the file containing the record was a sequential file. If, however, the NXR bit is specified, an alternative success code indicating successful access of a non-existent record (RMS\$_OK_RNF) will be returned. Any other stream also attempting to access or put a new record to record fifteen will receive a record-locked error (RMS\$_RLK). If the same stream that locked the non-existent record then attempts to put a new record there, an alternate success indicating that the record was already locked (RMS\$_OK_ALK) will be returned.

CHAPTER 11

RECORD-PROCESSING MACRO INSTRUCTIONS

VAX-11 RMS provides record-processing macro instructions that you can use to perform operations on individual records within a file, as opposed to operations performed on an entire file. These macro instructions, therefore, deal with fields in the record access block (RAB). At run time, the code of these macro instructions becomes expanded. At this time, calls are made to corresponding VAX-11 RMS services. See Chapter 5 for a description of the effect these fields have on the record operation.

After you open a file for processing with a \$OPEN or \$CREATE macro instruction, you can perform operations on the records in the file.

In most cases, you use a record-processing macro instruction with parameters indicating the symbolic address of the associated RAB and the address of any optional error or success completion routine you may have provided. You can also use the macro instruction without parameters, but you must then create an argument list in your program to define the values for these addresses (see Section 8.1).

Table 3-2 summarizes all the run-time processing macro instructions. This chapter deals only with the following macro instructions for record processing:

- Record access and current record context
 - \$GET
 - \$DELETE
 - \$PUT
 - \$FIND
 - \$UPDATE
- Record streams
 - \$CONNECT
 - \$DISCONNECT
- Synchronization with asynchronous operations
 - \$WAIT
- Miscellaneous operations
 - \$FLUSH
 - \$RELEASE
 - \$FREE
 - \$REWIND
 - \$NXTVOL
 - \$TRUNCATE

RECORD-PROCESSING MACRO INSTRUCTIONS

Chapter 10 discusses some of the general concepts involved in performing record operations. This chapter presents the details of the particular macro instructions that perform record operations, in alphabetical order.

\$CONNECT

11.1 ESTABLISHING A RECORD STREAM

The \$CONNECT macro instruction invokes the Connect service, which establishes a record stream by associating and connecting a RAB with a FAB. For sequential files, only one RAB can be connected to a FAB. For relative or indexed files, any number of RABs can be connected to a FAB, if the MSE bit was set in the file-sharing field (SHR) of the FAB when the file was opened or created. Each RAB represents an independent record stream.

When you issue a \$CONNECT macro instruction, VAX-11 RMS allocates an internal counterpart for the RAB. This counterpart consists of the necessary internal controls needed to support the stream, such as record pointers and request status information. All required I/O buffers are also allocated at this time. \$CONNECT also initializes the next record pointer to the first record. In indexed files, the key of reference establishes the index of the next record pointer.

You can issue a \$CONNECT macro instruction only to files that are already open.

Format

OPERATION	PARAMETERS
label: \$CONNECT	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$CONNECT macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

RECORD-PROCESSING MACRO INSTRUCTIONS

Table 11-1 lists the RAB fields that the connect service uses for input and output.

Table 11-1
Connect RAB Fields

Usage	Field Name	Description
Input	FAB	File access block address (used to access only the internal file identifier field of the FAB)
	KRF	Key of reference (used only with indexed files)
	MBC	Multiblock count (sequential disk files only)
	MBF	Multibuffer count
	ROP	Record-processing options (ASY, BIO, EOF, RAH, and WBH only)
Output	ISI	Internal stream identifier
	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Connect service are listed below.

Success:

RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
-----------	-----------------------------------

\$DELETE

11.2 DELETING A RECORD

The \$DELETE macro instruction invokes the Delete service, which removes an existing record from a relative or indexed file (you cannot use this macro instruction with sequential files). A record delete operation always applies to the current record. Therefore, immediately before you issue the \$DELETE macro instruction, you must lock the record by issuing a \$FIND or \$GET macro instruction.

RECORD-PROCESSING MACRO INSTRUCTIONS

Format

OPERATION	PARAMETERS
label: \$DELETE	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$DELETE macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-2 lists the RAB fields that the delete service uses for input and output. The VAX-11 RMS completion status codes are listed in Appendix A. However, to help anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the \$DELETE macro instruction are listed below.

Success:

RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_RNL	Warning; record not locked
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_WLK	Device write-locked

RECORD-PROCESSING MACRO INSTRUCTIONS

Table 11-2
Delete RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	ROP	Record-processing options (ASY and FDL bits only)
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value

\$DISCONNECT

11.3 TERMINATING A RECORD STREAM

The \$DISCONNECT macro instruction invokes the Disconnect service, which breaks the connection between a RAB and a FAB, thereby terminating a record stream. All system resources, such as I/O buffers and data structure space, are deallocated.

The close service (see Section 9.1) performs an implied disconnect for all record streams connected to the FAB.

Format

OPERATION	PARAMETERS
label: \$DISCONNECT	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$DISCONNECT macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

RECORD-PROCESSING MACRO INSTRUCTIONS

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-3 lists the RAB fields that the Disconnect service uses for input and output.

Table 11-3
Disconnect RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	ROP	Record-processing options (ASY bit only)
Output	ISI	Internal stream identifier (zeroed)
	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Disconnect service are listed below.

Success:

RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_WLK	Device write-locked

\$FIND

11.4 LOCATING A RECORD

The \$FIND macro instruction invokes the find service, which locates a specified record in a file and returns its record's file address in the RFA field of the RAB. This applies to all file organizations.

The main uses of the find service are:

- Skipping records when you are using the sequential record access mode (by issuing successive requests for find operations)
- Locking, but not retrieving, a record, thereby establishing a current record for an operation with a \$UPDATE, \$DELETE, or \$TRUNCATE macro instruction
- Establishing a random accessed starting point in a file for subsequent sequential access
- Randomly accessing records for Delete or Update operations, without modifying the next record context of sequential operations on the same stream. For a discussion of next record contexts, see Chapter 10.

Format

OPERATION	PARAMETERS
label: \$FIND	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$FIND macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-4 lists the RAB fields that the Find service uses for input and output.

RECORD-PROCESSING MACRO INSTRUCTIONS

Table 11-4
Find RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	KBF	Key buffer address (used only if RAC=KEY or if RAC=SEQ and the LIM option is selected in the ROP)
	KRF	Key of reference (used only with indexed files and if RAC=KEY)
	KSZ	Key size (used only if RAC=KEY or if RAC=SEQ and the LIM option is selected in the ROP)
	PBF	Prompt buffer address; applies to terminals only
	PSZ	Prompt buffer size? applies to terminals only
	RAC	Record access
	RFA	Record's file address (used only if RAC=RFA)
	ROP	Record-processing options (see sec.5.2.15)
	TMO	Time-out period
Output	BKT	Bucket code; set to the relative record number for relative files accessed sequentially
	RFA	Record's file address
	STS	Completion status code (also returned in Register 0)
	STV	Status value

The record address (RBF) and record size (RSZ) fields are undefined after a find service.

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Find service are listed below.

Success:

RMS\$_CONTROLC	Operation completed under Control C (terminals only)	C
RMS\$_CONTROLY	Operation completed under Control Y (terminals only)	Y
RMS\$_NORMAL	Operation successful	
RMS\$_OK_ALK	Record already locked	
RMS\$_OK_DEL	Deleted record accessed correctly	
RMS\$_OK_LIM	Record retrieved exceeds specified key value.	

RECORD-PROCESSING MACRO INSTRUCTIONS

RMS\$_OK_RLK	Record locked but read anyway
RMS\$_OK_RNF	Nonexistent record accessed correctly
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_BES	Warning; bad escape sequence (terminals only)
RMS\$_DEL	Record accessed by the RFA record access mode has been deleted
RMS\$_DNR	Device not ready
RMS\$_EOF	End of file
RMS\$_PES	Warning; partial escape sequence (terminals only)
RMS\$_RLK	Record locked by another stream
RMS\$_RNF	Record not found
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_TMO	Warning; time-out period expired (terminals only)
RMS\$_WLK	Device write-locked

\$FLUSH

11.5 WRITING OUT MODIFIED I/O BUFFERS

The \$FLUSH macro instruction invokes the Flush service, which writes out all modified I/O buffers and file attributes associated with the file. This ensures that all record activity up to the point at which this macro instruction executes is actually reflected in the file.

The Flush service is not required at any time. Even in the case of a \$CLOSE macro instruction, the Flush service is not needed, because the Close service implicitly performs the flush functions.

During asynchronous operations, you must wait for the completion of any I/O activity before issuing a \$FLUSH macro instruction. You may also issue a \$FLUSH macro instruction after having received notification of completion through an asynchronous system trap (AST).

RECORD-PROCESSING MACRO INSTRUCTIONS

Format

OPERATION	PARAMETERS
label: \$FLUSH	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$FLUSH macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-5 lists the RAB fields that the Flush service uses for input and output.

Table 11-5
Flush RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	ROP	Record-processing options (ASY bit only)
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes categorized as severe errors are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Flush service are listed below.

RECORD-PROCESSING MACRO INSTRUCTIONS

Success:

RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet completed

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_RSA	Record stream still active (asynchronous operations)

\$FREE

11.6 UNLOCKING ALL RECORDS

The \$FREE macro instruction invokes the free service, which unlocks all records that were previously locked for the record stream (see also the \$RELEASE macro instruction, Section 11.10). If no records are locked for the record stream, VAX-11 RMS returns a status code of RMS\$_RNL.

Section 10.5 describes the record-locking action.

Format

OPERATION	PARAMETERS
label: \$FREE	RAB=rab-address ERR=entry SUC=entry

label

A user-defined symbolic address for the \$FREE macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

RECORD-PROCESSING MACRO INSTRUCTIONS

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-6 lists the RAB fields that the Free service uses for input and output.

Table 11-6
Free RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, any error or warning completion status codes that can cause a failure for the Free service are listed below.

Success:

RMS\$_NORMAL Operation successful
RMS\$_PENDING Asynchronous operation not yet complete

Failure:

RMS\$_ACT File activity precludes operation
RMS\$_RNL Record not locked
RMS\$_RSA Record stream still active (asynchronous operations)

\$GET

11.7 RETRIEVING A RECORD

The \$GET macro instruction invokes the Get service, which causes a record to be retrieved from a file. The Get service is performed by using one of three possible record access modes, as specified by the record access field (RAC). The three modes are sequential (default), random by key, and random by record's file address (RFA access).

The sequential access mode is relevant for all file organizations as well as all devices. It is the only access mode allowed for non-disk devices, such as terminals, mailboxes, and magnetic tape devices. In this mode, records are retrieved from a given file in the same order that they were written to that file. This is not the case, however,

RECORD-PROCESSING MACRO INSTRUCTIONS

for records retrieved from indexed sequential files. Sequential Get operations on an indexed sequential file will return records by ascending key value. The next record's key of reference for sequential operations on indexed sequential files is established by one of the following operations:

1. \$ CONNECT
2. \$ REWIND
3. \$ FIND or \$GET using random by key access
4. \$ FIND or \$GET using random by RFA access.

In the case of the \$CONNECT, \$REWIND, and \$FIND or \$GET (using random by key access) operations, the key of reference is established by the key of reference field (KRF). The \$FIND or \$GET (using random by RFA access) operation, on the other hand, always sets the key of reference to the primary key.

Random by key record access mode is used to retrieve records by key value. For all relative files, and for sequential files with fixed-length records, the key value is the relative record number. For indexed sequential files, the key value is dependent on the data type of the specified key of reference. The key value is used to search the index of the specified key of reference by value in order to locate the desired record. A random access by key also establishes the next record for subsequent sequential retrieval. It may be used in this way to establish a starting point for sequential retrieval of records at other than the beginning of the file.

Random by record's file address (RFA) access is used to randomly retrieve records from disk files. In order to determine a record's RFA, the record must have been previously accessed. The RFA is output from Find, Get, and Put operations.

If when opening the file, only sequential operations (the SQO option in the FOP field of the FAB) are specified, then the random access of records in that file is not permitted. A further discussion of the different record access modes, as well other record-processing concepts, is contained in Chapter 10 and in Appendix B.

VAX-11 RMS uses the standard terminator set when performing input operations from terminal devices. The terminating character is returned in the low order word of the status value field (STV0). See the chapter on the terminal driver contained in the VAX/VMS I/O User's Guide. In addition to terminating the read request, the <CTRL/Z> character is treated as an end-of-file marker by VAX-11 RMS. If you enter a response to a read request, VAX-11 RMS will return the completion status end-of-file (RMS\$_EOF). Data entered prior to the <CTRL/Z> will be returned successfully. The next get request will return a single end-of-file error (RMS\$_EOF) without accepting any further input from the device. A subsequent get request, however, will resume the acceptance of input from the device.

VAX-11 RMS also supports the use of escape sequences from terminal devices that are accessed locally and have escape sequence enabled. Escape sequences for a terminal are enabled by the SET TERMINAL command described in the VAX/VMS Command Language User's Guide. Escape sequences are returned in the record buffer. The record size (RSZ) is the offset (within the buffer (RBF)) to the beginning of the escape sequence. The high order word of the status value field (STV2) will contain the length of the escape sequence. When a partial escape sequence error (RMS\$_PES) is returned, the remaining characters in the

RECORD-PROCESSING MACRO INSTRUCTIONS

escape sequence will be returned by the next read request from the terminal.

Mailboxes may be used to synchronize activity across cooperating processes. Normally, a Get service from a mailbox device will not be completed until a record is present in the mailbox. When the Get service is completed, the status value field (STV) will contain the process identification (PID) of the process that put the record into the mailbox. However, if the time out (TMO) record option (ROP) is specified with a value of zero in the time-out field, and if no messages are present in the mailbox, then the Get operation will return an end-of-file error (RMS\$ EOF). This technique assures your process of an immediate return, whether or not messages are present in the mailbox.

The STV field contains additional status information for a number of situations. When the completion status is a record-too-big warning (RMS\$ RTB), the STV contains the total record size. When the device is record oriented (e.g., terminals and mailboxes), the second longword of the I/O status block is returned in the STV field, whenever the completion status (STS) is a success code. The alternate field definitions of RAB\$W_STV0 and RAB\$W_STV2 are provided to reference the respective low and high order words of the status value field. The record size field (RSZ) always reports the amount of data returned, regardless of the completion status (STS). The presence of valid data on error conditions may then be detected by checking the record size field.

The Get service always requires the presence of a user record area, as specified by the user record area address (UBF) and user record area size (USZ) fields in the RAB.

Format

OPERATION	PARAMETERS
label: \$GET	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$GET macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

RECORD-PROCESSING MACRO INSTRUCTIONS

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-7 lists the RAB fields that the Get service uses for input and output.

Table 11-7
Get RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	KBF	Key buffer address (used only if RAC=KEY or if RAC=SEQ and the LIM option is selected in the ROP)
	KRF	Key of reference (used only with indexed files and if RAC=KEY)
	KSZ	Key buffer size (used only if RAC=KEY or if RAC=SEQ and the LIM option is selected in the ROP)
	PBF	Prompt buffer address; applies to terminals only
	PSZ	Prompt buffer size; applies to terminals only
	RAC	Record access mode
	RFA	Record's address (used only if RAC=RFA)
	RHB	Record header buffer; used for variable with fixed control records
	ROP	Record-processing options (see Section 5.2.15)
	TMO	Time-out period; applies to terminals and mailboxes only
	UBF	User record area address
	USZ	User record area size
Output	BKT	Bucket code; set to the relative record number for relative files when the record access mode is sequential
	RBF	Record address
	RFA	Record's file address
	RSZ	Record size
	STS	Completion status code (also returned in Register 0)
	STV	Status value (contains a terminator character for terminal input or the record length if the requested record is too large for the user buffer area)

RECORD-PROCESSING MACRO INSTRUCTIONS

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Get service are listed below.

Success:

RMS\$_CONTROL C	Operation completed under Control C (terminals only)
RMS\$_CONTROL Y	Operation completed under Control Y (terminals only)
RMS\$_NORMAL	Operation successful
RMS\$_OK_ALK	Record already locked
RMS\$_OK_DEL	Deleted record accessed correctly
RMS\$_OK_LIM	Retrieved record exceeds specified key value
RMS\$_OK_RLK	Record locked but read anyway
RMS\$_OK_RNF	Nonexistent record accessed correctly
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_BES	Warning; bad escape sequence (terminals only)
RMS\$_DEL	Record accessed by the RFA record access mode has been deleted
RMS\$_DNR	Device not ready
RMS\$_EOF	End of file
RMS\$_PES	Warning; partial escape sequence (terminals only)
RMS\$_RLK	Record locked by another stream
RMS\$_RNF	Record not found
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_RTB	Warning; record too large for user buffer
RMS\$_TMO	Warning; time-out period expired
RMS\$_TNS	Warning; terminator not seen (terminals only)
RMS\$_WLK	Device write-locked

\$NXTVOL

11.8 CONTINUE PROCESSING ON NEXT VOLUME

The \$NXTVOL macro instruction invokes the Next volume service. This service applies only to files on magnetic tape volumes. Use this macro instruction when you want to proceed to the next volume in the set before the end of the current volume (EOV label) is reached on input, or before the end of tape (EOT mark) is reached on output. VAX-11 RMS will then position to the first file section on the next volume. File sections occur when a file is written on more than one volume, the portion of the file on each of the volumes constituting a file section.

For input files, the following occurs:

- If the current volume is the last volume of the set, VAX-11 RMS reports end-of-file.
- If another file section exists, the next volume is mounted. When necessary, the current volume is rewound and a request to mount the next volume is issued to the operator.
- The header label (HDR1) of the file section on the newly mounted volume is read. If it is not the volume being sought, the operator is requested to mount the correct volume.

For output files, the following occurs:

- The file section on the current volume is closed with the appropriate end-of-volume labels, and the volume is rewound.
- The next volume is mounted.
- A file with the same file name and the next higher file section number is opened for output, and processing continues.

If operating asynchronously, you must wait for the completion of any I/O activity on this volume before issuing a \$NXTVOL macro instruction.

The Next volume service performs a flush operation for write-accessed volumes (see Section 11.5), thus writing the I/O buffers on the current volume before creating the next file section. If this is an input-only file, all records currently contained in the I/O buffers are lost, and the next get call will return the first record on the next volume.

RECORD-PROCESSING MACRO INSTRUCTIONS

Format

OPERATION	PARAMETERS
label: \$NXTVOL	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$NXTVOL instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-8 lists the RAB fields that the Next volume service uses for input and output.

Table 11-8
Next Volume RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	ROP	Record-processing options (ASY bit only)
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Next volume service are listed below.

RECORD-PROCESSING MACRO INSTRUCTIONS

Success:

RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_DPE	Device positioning error
RMS\$_RSA	Record stream still active (asynchronous operations)

\$PUT

11.9 WRITING A RECORD TO A FILE

The \$PUT macro instruction invokes the Put service, which inserts a record into a file. The new records can be placed either at the end of the file (sequential and relative file organizations) or in empty record cells in place of deleted records (relative file organization). Location of new records in an indexed file is controlled by VAX-11 RMS, which examines the contents of the primary key field of the record to determine where to write the record into the file.

When using sequential files with the sequential record access mode, you normally write records only at the end of the file. These records cannot have a length greater than the maximum you specified when you created the file. You can use random by key record access mode to write fixed-length records in a sequentially organized disk file. The Truncate on put (TPT) record option (ROP) may be specified for sequential files. This option allows new records to be inserted in existing files, in positions of the file other than at the end of file. When the file is closed, it will automatically be truncated to the new end of file. Truncate access is required to use this option. When the file is closed, it will automatically be truncated to a new end of file, immediately after the last record is inserted (whether or not data had existed after that point).

In a relative file, you can use either sequential or random by key record access mode. Records cannot be larger than the size specified at file creation, and the record's relative record number must not exceed the maximum record number established for the file. Normally, if the target record cell for a Put operation contains a record, a record-already-exists error (RMS\$_REX) will be returned as the completion status (STS). If you specify the Update-if (UIF) record option (ROP), however, VAX-11 RMS will overwrite the existing record, instead of returning an error message.

In an indexed file, you can use either sequential or random by key record access mode. When sequential access is used to put (insert) records, the primary key value of the record to be put must be equal to or greater than the primary key value of the preceding record. The records cannot be larger than the size established (if a maximum length was specified) when the file was created. Each record written must contain a complete primary key, but the records do not have to contain all alternate keys. If alternate keys are partially or

RECORD-PROCESSING MACRO INSTRUCTIONS

completely missing because of record length, VAX-11 RMS will not make an entry for that new record in the associated alternate index(es). Put operations to an indexed file do not require a key value or key of reference. VAX-11 RMS determines where to write the record by examining the contents of the primary key in the record.

When inserting records into an indexed sequential file, VAX-11 RMS compares the key values (primary and alternate) in the record with the key values of records already existing in the file. This comparison determines if the writing of the record would result in the presence of duplicate key values among records of the file. If duplicates would occur, VAX-11 RMS verifies that duplicates are allowed. If duplicates are not allowed for a particular key, VAX-11 RMS rejects the operation with an RMS\$_DUP error code. However, if duplicates are allowed, VAX-11 RMS performs the operation. Subsequent sequential operations on a given index will always retrieve records with identical key values in the order in which the records were inserted.

If you specify the Update-if (UIF) record option (ROP) when duplicates are not allowed on the primary key, VAX-11 RMS will simply overwrite an existing record with the same primary key value, rather than returning a duplicate record error (RMS\$_DUP). Alternate key values will be modified to reflect the newly-inserted record. It will appear as if an Update operation is being performed on the existing record. When the Update-if option is used, update access to the file is required. If update access is not permitted for the file, then a Put operation (that becomes an Update operation when this option is selected) will fail, and a file-access error (RMS\$_FAC) will be returned.

Mailboxes may be used to synchronize activity across cooperating processes. Normally, a Put service to a mailbox will not be completed until another accessor to the mailbox reads the record. When the Put service is completed, the status-value field (STV) will contain the Process identification (PID) of the process that read the record. If the Time-out (TMO) record option (ROP) is specified with a time-out period of zero, the Put service will not wait for another accessor to read the record.

The record address (RBF) and record size (RSZ) are required for all Put operations. Some options for the Put operations, however, require additional fields. Following the completion of a successful Put service, the record's file address (RFA) is always returned. A description of the concepts relevant to understanding the interrelationship between the Put operation and other record operations is contained in Chapter 10.

Format

OPERATION	PARAMETERS
label: \$PUT	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$PUT macro instruction; optional.

RECORD-PROCESSING MACRO INSTRUCTIONS

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

Suc=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-9 lists the RAB fields that the Put service uses for input and output.

Table 11-9
Put RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	KBF	Key buffer address (used only if RAC=KEY and the file is a relative file)
	KSZ	Key size (used only if RAC=KEY and the file is a relative file)
	RAC	Record access mode
	RBF	Record address
	RHB	Record header buffer; only applies to variable with fixed control records
	RSZ	Record size
	ROP	Record-processing options (ASY, CCO, RLK, TPT, TMO, UIF, ULK and WBH only)
	TMO	Time-out; mailboxes only. ¹
Output	BKT	Bucket code; set to the relative record number for sequential access to relative files
	RFA	Record's file address
	STS	Completion status code (also returned in Register 0)
	STV	Status value

¹On the successful completion of a put service to a record-oriented device, the STV field will contain the second longword of the I/O status block. See the VAX/VMS I/O User's Guide for details on specific devices.

RECORD-PROCESSING MACRO INSTRUCTIONS

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Put service are listed below.

Success:

RMS\$_CONTROLC	Operation completed under Control C (terminals only)
RMS\$_CONTROLO	Operation completed under Control O (terminals only)
RMS\$_CONTROLY	Operation completed under Control Y (terminals only)
RMS\$_NORMAL	Operation successful
RMS\$_OK_ALK	Record already locked
RMS\$_OK_DUP	Record inserted has duplicate key value
RMS\$_OK_IDX	Record successfully inserted, but error occurred on index update which could cause slow access
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_DUP	Duplicate key detected (see above text)
RMS\$_EXT	File extend error
RMS\$_PRV	Privilege violation; access denied
RMS\$_REX	Record already exists in target record cell
RMS\$_RLK	Record locked by another task
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_RVU	Error updating RRVs
RMS\$_WLK	Device write-locked

\$RELEASE

11.10 UNLOCKING A RECORD

The \$RELEASE macro instruction invokes the Release service, which unlocks the record pointed to by the contents of the record's file address RFA field of the RAB (see also the \$FREE macro instruction, Section 11.6). If the named record is not locked, VAX-11 RMS returns a status code of RMS\$_RNL.

RECORD-PROCESSING MACRO INSTRUCTIONS

Section 10.5 describes record locking.

Format

OPERATION	PARAMETERS
label: \$RELEASE	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$RELEASE macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-10 lists the RAB fields that the Release service uses for input and output.

Table 11-10
Release RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	RFA	Record's file address
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value

RECORD-PROCESSING MACRO INSTRUCTIONS

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Release service are listed below.

Success:

RMS\$_NORMAL Operation successful
RMS\$_PENDING Asynchronous operation not yet complete

Failure:

RMS\$_ACT File activity precludes operation
RMS\$_RNL Warning; record not locked
RMS\$_RSA Record stream still active (asynchronous operations)

\$REWIND

11.11 POSITIONING TO THE FIRST RECORD

The \$REWIND macro instruction invokes the Rewind service, which sets the current context of a stream to the first record in the file. VAX-11 RMS alters the context of the next record to indicate the first record as being the next record. The Rewind service implicitly performs the Flush and Free services, writing out all I/O buffers and releasing all locked records. This service is valid for all file organizations on disk volumes and for sequential files on tape volumes. For indexed files, the KRF field establishes the index to be used for subsequent sequential accesses. You cannot rewind a unit record device (card reader or printer) or a terminal.

Format

OPERATION	PARAMETERS
label: \$REWIND	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$REWIND macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

RECORD-PROCESSING MACRO INSTRUCTIONS

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-11 lists the RAB fields that the Rewind service uses for input and output.

Table 11-11
Rewind RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	KRF	Key of reference (used only with indexed files)
	ROP	Record-processing options (ASY bit only)
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure are listed below.

Success:

RMS\$_NORMAL Operation successful
RMS\$_PENDING Asynchronous operation not yet complete

Failure:

RMS\$_ACT File activity precludes operation
RMS\$_BOF Warning; file is already at beginning of file
RMS\$_DNR Device not ready
RMS\$_DPE Device positioning error
RMS\$_EOF End of file
RMS\$_RSA Record stream still active (asynchronous operations)
RMS\$_WLK Device write-locked

\$TRUNCATE

11.12 TRUNCATING A SEQUENTIAL FILE

The \$TRUNCATE macro instruction invokes the Truncate service, which truncates records from the end of a sequential file. Note that you can only truncate a sequential file (you cannot use this service for a relative or indexed file) and the file must be open for exclusive access (the file-sharing field of the FAB must be set or defaulted to NIL). The file access field (FAC) must specify truncate access (TRN). The Truncate service deletes the record indicated as the Current Record, and all following records. You can only use this service immediately after successful execution of a \$GET, \$FIND, or \$UPDATE macro instruction (thereby setting the context of the Current Record).

VAX-11 RMS declares an end-of-file at the starting record position for the truncation, and then causes the context of the Next record to be set to this end of file. You can then add records to the file by issuing successive \$PUT macro instructions.

Format

OPERATION	PARAMETERS
label: \$TRUNCATE	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$TRUNCATE macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-12 lists the RAB fields that the truncate service uses for input and output.

RECORD-PROCESSING MACRO INSTRUCTIONS

Table 11-12
Truncate RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	ROP	Record-processing options (ASY only)
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Truncate service are listed below.

Success:

RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_DPE	Device positioning error
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_WLK	Device write-locked

\$UPDATE

11.13 UPDATING AN EXISTING RECORD

The \$UPDATE macro instruction invokes the Update service, which modifies (updates) the contents of an existing record in a disk file only. The record to be updated must first be locked by this stream, either by a \$FIND or \$GET macro instruction. You cannot use locate mode; you must supply a buffer.

For sequential files, the record length cannot change. For relative files with variable-length or variable with fixed-length control records, the length of the replacement record can differ from the length of the original record, but cannot be larger than the maximum size you set when you created the file.

RECORD-PROCESSING MACRO INSTRUCTIONS

For indexed files, the length of the replacement (updated) record written by the \$UPDATE macro instruction may be different from the original record; restrictions, however, apply to the replacement record in an indexed file:

- The length of the replacement record cannot exceed the maximum size defined at file creation.
- Each replacement record must be large enough to contain a complete primary key, but the replacement record does not have to contain all alternate keys. If an alternate key is partially or completely missing in the replacement record, the key must have the characteristic that the values can change; this is true also if the replacement record contains a key that was not present in the original record.

Update operations to an indexed file do not require a key value or key of reference. Before writing the record, VAX-11 RMS compares the key values (primary and alternate) in the replacement record with the key values of original record already existing in the file. This comparison takes into account the defined characteristics of each key. For example, if a particular key is not allowed to change, VAX-11 RMS rejects the operation with an RMS\$CHG error code if the replacement record contains an altered value in the associated key. Similarly, this comparison determines if the replacement record would result in the presence of duplicate key values among records of the file. If duplicates would occur, VAX-11 RMS verifies the defined characteristics for the keys being duplicated. If duplicates are not allowed for a particular key, VAX-11 RMS rejects the operation with an RMS\$DUP error code. However, if duplicates are allowed, VAX-11 RMS performs the operation.

Subsequent sequential operations on a given index will always retrieve records with identical key values in the order in which the records were written.

Format

OPERATION	PARAMETERS
label: \$UPDATE	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$UPDATE macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

RECORD-PROCESSING MACRO INSTRUCTIONS

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 11-13 lists the RAB fields that the update service uses for input and output.

Table 11-13
Update RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	RBF	Record address
	RHB	Record header buffer; applies only to variable with fixed control records
	ROP	Record-processing options (ASY and WBH only)
	RSZ	Record size
Output	RFA	Record's file address
	STS	Completion status code (also returned in Register 0)
	STV	Status value

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Update service are listed below.

Success:

RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete
RMS\$_OK_DUP	Duplicate key detected
RMS\$_OK_IDX	Record was inserted, but error occurred on index update which could cause slow access

RECORD-PROCESSING MACRO INSTRUCTIONS

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_RNL	Warning; record not locked
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_WLK	Device write-locked

\$WAIT

11.14 STALL FOR I/O COMPLETION

The \$WAIT macro instruction invokes the wait service, which determines when an asynchronous record operation completes. Upon completion of the operation, VAX-11 RMS returns control to your program at the point following the \$WAIT macro instruction. Any completion routines specified on the operation being awaited are also executed before VAX-11 RMS returns control (unless ASTs are disabled).

The \$WAIT macro instruction takes no parameters to define entry points for user-written completion routines; the completion routines are specified by the operation being awaited.

Format

OPERATION	PARAMETERS
label: \$WAIT	RAB=rab-address

label

A symbolic address for the \$WAIT macro instruction; optional.

RAB=rab-address

This parameter defines the symbolic address of the RAB; either the RAB whose I/O request is in progress, or some other RAB.

Table 11-14 lists the RAB fields that the Wait service uses for input and output.

RECORD-PROCESSING MACRO INSTRUCTIONS

Table 11-14
Wait RAB Fields

Usage	Field Name	Description
Input	ISI	Internal stream identifier
	STS	Completion status code
Output	STS	Completion status code (also returned in Register 0)

The VAX-11 RMS completion status codes for the wait service are determined by the operation being awaited, unless the address of the RAB specified for the wait is not the same as that specified for the awaited operation. In this case, RMS\$_NORMAL is returned.

CHAPTER 12

PERFORMING BLOCK I/O

Besides the record access provided by the sequential, random by key, and random by record's file address record access modes, VAX-11 RMS also provides another means to access data in a file: block I/O.

Block I/O operations let you directly read or write the blocks of a file. These operations are provided for users who must keep system overhead to a minimum and need no interpretation of file data as logical records, yet still want to take advantage of the easy file access of VAX-11 RMS. Block I/O is an intermediate step between the VAX-11 RMS record operations and direct use of VAX/VMS input/output system services.

You specify Block I/O for a record stream by setting the BIO bit in the file access field of the file access block (FAB; see Section 4.2.10) as input to the \$OPEN or \$CREATE macro instructions. If you intend to write to the file, you must set the PUT bit in the file access field. If you want to read from the file, you must set the GET bit in the file access field.

You cannot perform Block I/O operations on files on which record operations are already being performed. Conversely, you cannot perform record operations on files on which Block I/O operations are being performed. However, VAX-11 RMS allows you to set the BRO bit in the file access field of the FAB, indicating that operations can switch from Block I/O to record operation and vice versa when an operation is completed (but not using both at the same time). Only the sequential file organization allows this switching. For other file organizations, setting of the BRO bit of the file access field merely allows the decision about performing block or record operations to be delayed until the first RAB is connected. If the BIO bit is set in the record options field of the RAB, only Block I/O operations will be permitted; if the BIO bit is clear, only record operations will be permitted. All connected record streams must be connected in the same manner; that is, there can be no mixing of Block and Record I/O.

If you do mix modes of operation for sequential files, you must exercise caution, as the context of the current record, next record, and the next block pointer (see NOTES below) are all undefined when you switch operations on disk devices. Therefore, the operation that initiates the switch must not use sequential record access mode. For magnetic tape devices, the context of the next record or next block indicates the start of the following Block on the tape for the operation initiating the switch.

PERFORMING BLOCK I/O

NOTES

1. If you set the BRO bit in the file access field of the FAB for the sequential file organization, you indicate that you want to mix block I/O and record operations. If, once the file is open, you want only to perform block I/O, you can set the BIO bit in the record-processing options field of the RAB. This operation overrides the setting of the BRO bit for this record stream, and acts as a flag to the \$CONNECT macro instruction, indicating that no VAX-11 RMS I/O buffers need be allocated (but you must still allocate buffers in the user program for block I/O operations).
2. If you set the BRO bit when creating an indexed file, the key definition XABs for that file must be present. For a create service to the relative or indexed file organization, specifying the BIO bit in the file access field of the FAB causes VAX-11 RMS to omit prologue processing and initial space pre-zeroing in relative files. Allocated space pre-zeroing is also omitted for the extend service when connected for block I/O.
3. For files of unknown organization or undefined record format, block I/O is the only form of processing allowed. Processing proceeds identically to that for block I/O to the relative file organization.

Three macro instructions are provided for performing block I/O.

- \$READ -- transfers a specified number of bytes into memory
- \$SPACE -- positions a file forward or backward a specified number of blocks
- \$WRITE -- writes a specified number of bytes to a file

In addition, you can use the following macro instructions on a record stream connected for Block I/O operations:

- \$DISCONNECT
- \$NXTVOL
- \$FLUSH
- \$REWIND

These instructions, which are described in Chapter 11, perform miscellaneous operations or disconnect the record stream, and do not work on the contents of the records themselves.

PERFORMING BLOCK I/O

For sequential Block I/O operations to disk files, VAX-11 RMS maintains an internal next block pointer (NBP) that:

- Points to the beginning of the file after execution of a \$CONNECT macro instruction if the EOF bit is cleared in the record-processing options field of the record access block (RAB), or if the EOF bit is set, NBP points to the block following the end of file. For indexed files, setting EOF is not permitted.
- Points to the block following the highest numbered block transferred by a read or write service (\$READ or \$WRITE macro instructions).
- Points to the next block after an operation with the \$SPACE macro instruction.

The \$BLOCK I/O macro instructions deal with fields in the RAB; Chapter 5 describes the effect of these fields on the operations.

You indicate the symbolic address of the associated RAB through a parameter on each Block I/O macro instruction you are using, and the address of any optional error or success completion routine you may have provided. However, you can also use the macro instruction without parameters, but you must then create an argument list in your program to define the values for these addresses (see Section 8.1).

\$READ

12.1 TRANSFER TO MEMORY

The \$READ macro instruction invokes the Read service, which retrieves a specified number of bytes from a file (on a block boundary) and transfers them to memory. A Read operation using block I/O can be performed on any file organization.

To use this macro instruction, you must:

1. Supply a buffer area into which VAX-11 RMS is to transfer data (user record area address field).
2. Indicate the number of bytes to be transferred (user record area size field).
3. Indicate the first virtual block number (VBN) for the transfer (bucket number field). If the value for the VBN is zero, the transfer will start with the block indicated by the NBP.

Format

OPERATION	PARAMETERS
label: \$READ	RAB=rab-address ERR=entry SUC=entry

PERFORMING BLOCK I/O

label

A symbolic address for the \$READ macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 12-1 lists the RAB fields that the Read service uses for block I/O.

Table 12-1
Read RAB Fields

Usage	Field Name	Description
Input	BKT	Bucket number; must contain the virtual block number of the first block to be read
	ISI	Internal stream identifier
	ROP	Record-processing options; (ASY bit only)
	UBF	User record area address
	USZ	User record area size; indicates the length of transfer, in bytes
Output	RBF	Record address
	RFA	Record's file address; contains the virtual block number of the first block transferred
	RSZ	Record size; indicates the actual number of bytes transferred
	STS	Completion status code (also returned in Register 0)
	STV	Status value

PERFORMING BLOCK I/O

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the \$READ macro instruction are listed below.

Success:

RMS\$_CONTROLC	Operation completed under Control C
RMS\$_CONTROLY	Operation completed under Control Y
RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_EOF	End of file; checking for the logical end of file is performed for the sequential file organization only. If an end-of-file error occurs, it implies that the first virtual block number specified was at or past the end of the file. If the end-of-file pointer occurs during a transfer, the record size field is set to the number of bytes before the logical end of file. For the relative file organization, this status code indicates an attempt to read past the end of the currently allocated space.
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_TMO	Warning; time-out period expired
RMS\$_WLK	Device write-locked

\$SPACE

12.2 POSITIONING TO A BLOCK

The \$SPACE macro instruction invokes the Space service, which lets you position a file forward or backward a specified number of blocks.

This macro instruction is intended primarily for use with magnetic tape files; the tape is spaced the number of blocks specified in the bucket number field. If the value in this field is positive, the tape spaces forward; if the value is negative, the tape spaces backward. For disk files, the NBP is updated to reflect the new sequential operation position.

PERFORMING BLOCK I/O

Format

OPERATION	PARAMETERS
label: \$SPACE	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$SPACE macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written success completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 12-2 lists the RAB fields that the Space service uses as input and output.

Table 12-2
Space RAB Fields

Usage	Field Name	Description
Input	BKT	Bucket number; indicates the number of blocks to space forward (positive value) or backward (negative value)
	ISI	Internal stream identifier
	ROP	Record-processing options; ASY bit only
Output	STS	Completion status code (also returned in Register 0)
	STV	Status value (set to number of blocks actually spaced; positive value always)

PERFORMING BLOCK I/O

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Space service are listed below.

Success:

RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_BOF	File is at beginning-of-file (backspace operations)
RMS\$_DNR	Device not ready
RMS\$_DPE	Device positioning error
RMS\$_EOF	File is at end-of-file (forward space operations)
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_WLK	Device write-locked

\$WRITE

12.3 WRITE TO A FILE

The \$WRITE macro instruction invokes the Write service, which transfers a user-specified number of bytes, beginning on a block boundary, to a VAX-11 RMS file of any file organization.

You indicate the number of bytes to be written in the record size field of the RAB, and indicate the address of the buffer for the transfer in the record address field. In the bucket number field, you indicate the virtual block number of the first block to be written; if this number is 0, the transfer starts with the block indicated by the NBP.

For sequential files, the file is automatically extended if you write a block past the end of the currently allocated space. For relative and indexed sequential files, you must use the \$EXTEND macro instruction.

For sequential files, VAX-11 RMS maintains a logical end of file to correspond to the last block and highest byte written within the block.

PERFORMING BLOCK I/O

Format

OPERATION	PARAMETERS
label: \$WRITE	RAB=rab-address ERR=entry SUC=entry

label

A symbolic address for the \$WRITE macro instruction; optional.

RAB=rab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the RAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 12-3 lists the RAB fields that the Write service uses as input and output.

Table 12-3
Write RAB Fields

Usage	Field Name	Description
Input	BKT	Bucket number; must contain the virtual block number of the first block to be written
	ISI	Internal stream identifier
	RBF	Record address
	ROP	Record-processing options (ASY and TPT bits only)
	RSZ	Record size; indicates the transfer length, in bytes.
Output	RFA	Record's file address; contains the virtual block number of the first block transferred.
	STS	Completion status code (also returned in Register 0)
	STV	Status value; contains the actual number of bytes transferred if an end-of-file error occurs.

PERFORMING BLOCK I/O

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Write service are listed below.

Success:

RMS\$_CONTROLC	Operation completed under Control C
RMS\$_CONTROLO	Operation completed under Control O
RMS\$_CONTROLY	Operation completed under Control Y
RMS\$_NORMAL	Operation successful
RMS\$_PENDING	Asynchronous operation not yet complete.

Failure:

RMS\$_ACT	File activity precludes operation
RMS\$_DNR	Device not ready
RMS\$_EOF	End of file; for the sequential file organization, this error implies that the file could not be extended
RMS\$_EXT	ACP file extend error
RMS\$_PRV	File protection violation
RMS\$_RSA	Record stream still active (asynchronous operations)
RMS\$_WLK	Device write-locked

12.4 NON-FILE-STRUCTURED OPERATIONS

VAX-11 RMS lets you perform non-file-structured operations, that is, operations that deal with magnetic tape or disk volumes directly rather than through the Files-11 structure.

Non-file-structured operations also are known as Logical I/O operations. Logical I/O is the reading and writing of data in blocks. For magnetic tape, each block of the tape is read or written with no interpretation of labels. For disk, the starting block for a transfer is identified by a logical block number (LBN). Since LBNs are volume-relative (see Appendix B), no file-relative translation is required to determine the blocks to transfer.

You can perform non-file-structured operations under the following conditions.

1. For file devices that have been mounted as Files-11 volumes, you must set the NFS bit in the file-processing options field (FOP) of the FAB as input to the create or open service.
2. For file devices mounted as foreign (i.e., non-file-structured), VAX-11 RMS performs non-file-structured operations regardless of the state of the NFS bit.

PERFORMING BLOCK I/O

3. For nonfile devices, non-file-structured operations occur always.
4. If the NFS bit is set, the I/O channel is assigned in the mode of the caller, thus allowing I/O calls to be performed directly, if desired.
5. You must have the appropriate privileges to perform non-file-structured operations (logical I/O privilege) if the device is mounted as a Files-11 device.
6. Either block I/O or the get and put services are allowed. For magnetic tape, blocking information must be specified on the MOUNT command (see the VAX/VMS Command Language User's Guide), using the /RECORD qualifier; this allows the blocking and unblocking of fixed-length records, with records not crossing block boundaries. For disk, each block is read as a fixed-length record of 512 bytes.
7. The file specification needs only the device and unit number.

If the above conditions have been met, VAX-11 RMS will change its operations to include the following:

1. The Block I/O services including space are permitted, even if not in block I/O mode.
2. The Rewind service rewinds the entire magnetic tape.
3. If the Close service is performed to a write-accessed magnetic tape, and if the last operation performed was a Write operation, then two tape marks followed by a backspace will be output. This operation allows the creation of multiple files. On input, end-of-file errors cause the tape mark to be skipped.
4. For disk, the normal input of the bucket code field (BKT) of the RAB for Read and Write services specifies the logical block number (LBN) rather than the virtual block number (VBN). Since logical block numbers start at 0 and virtual block numbers start at 1, a problem may arise when you want to access LBN 0 (a 0 in the bucket code field indicating sequential operations). However, you can access LBN 0 by setting the bucket code field to 0 immediately after a Connect or Rewind service (or by issuing an appropriate Space service to backspace to the beginning of the volume).
5. For the Get and Put service, random access by key (RAC=KEY), set the key buffer pointed to by the key buffer address field to the starting LBN.

CHAPTER 13

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

VAX-11 RMS provides macro instructions that perform actions related to the file specification. These macro instructions are used only for relatively complex operations, such as wild card character processing, as their functions are normally performed by the Open and Create services.

These macro instructions, therefore, deal with fields in the file access block (FAB), and the name (NAM) block. Chapters 4 and 7 describe the effects of these fields for the FAB and NAM block, respectively. The file specification processing macro instructions are:

- \$ENTER
- \$PARSE
- \$REMOVE
- \$RENAME
- \$SEARCH

You indicate the symbolic address of the associated FAB through a parameter on the file specification processing macro instructions. You do not indicate the NAM block on the macro instructions; rather, you associate this NAM block with the FAB through the name block address field of the FAB.

On the file specification processing macro instructions, you can also use a parameter to indicate the address of any optional error or success completion routine you may have provided. You can use the macro instruction without parameters, but you must then create an argument list in your program to define the values for these addresses (see Section 8.1).

\$ENTER

13.1 ENTER A FILE NAME

The Enter service, which you invoke with the \$ENTER macro instruction, inserts a file name into a directory. This is performed automatically by the Create service (unless either the TMP or TMD bit is set in the file-processing options field of the FAB). The enter service, however, allows you to perform this step separately.

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

When you enter a file name into a directory, no file can already be open with the FAB, and no wild card character specifications are allowed.

The Enter service requires many NAM block fields as input. You normally precede the Enter service with an Open, Create, or Parse service (see Section 13.2), and a Search service (see Section 13.5), specifying the same FAB and NAM block for each service.

Format

OPERATION	PARAMETERS
label: \$ENTER	FAB=fab-address ERR=entry SUC=entry

label

A symbolic address for the \$ENTER macro instruction; optional.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 13-1 lists the fields in both the FAB and NAM block that the enter service uses as input and output.

The optional resultant string is moved to the buffer described by the resultant string area address (RSA) and size (RSS) fields of the NAM block (only if both these fields are nonzero).

If the file version number of the name string described by the expanded string length and area address fields of the NAM block is either not present or 0, the Enter service scans the entire directory. It assigns a version number one higher than the highest found (or 1 if none is found).

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

Table 13-1
Enter Fields

Usage	Control Block	Field Name	Description
Input	FAB	IFI	Internal file identifier (must be zero)
		NAM	Name block address
	NAM	DID	Directory identification; file name and identifier are entered into this directory
		DVI	Device identification of the device containing directory where file name is to be entered
		ESA	Expanded string area address; contains file name, type, and version to be entered
		ESL	Expanded string length
		FID	File identification of file to be entered into directory
		RSA	Resultant string area address
		RSS	Resultant string size
Output	FAB	STS	Completion status code (also returned in Register 0)
		STV	Status value
	NAM	RSL	Resultant string length

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Enter service are listed below.

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_DNF Directory not found
RMS\$_DNR Device not ready
RMS\$_ENT ACP enter function failed
RMS\$_FNF File not found
RMS\$_PRV File protection violation
RMS\$_WLK Device write-locked

\$PARSE

13.2 PARSE A FILE SPECIFICATION STRING

The \$PARSE macro instruction invokes the parse service, which analyzes the file specification string (as described in Section 8.2) and fills in various NAM block fields. The functions of the Parse service are performed automatically as part of the Open, Create, and Erase services.

When you parse a file name string, there must be no file already open in conjunction with the FAB. Section 8.2 describes the process of parsing a file specification. Appendix C describes the complete file specification syntax.

One function of the Parse service is to prepare the FAB and NAM blocks for wild card character processing to be used in the Search service. If wild card characters are present in the file specification, RMS allocates internal data structures to store the wild card character context for subsequent searches. This space is released when "\$SEARCH" encounters a No-More-Files condition (in which case an RMS\$ NMF ERROR is returned) or when another parse is performed using the same FAB and NAM blocks.

Format

OPERATION	PARAMETERS
label: \$PARSE	FAB=fab-address ERR=entry SUC=entry

label

A symbolic address for the \$PARSE macro instruction; optional.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 13-2 lists the fields in both the FAB and NAM block that the Parse service uses as input and output.

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

Table 13-2
Parse Fields

Usage	Control Block	Field Name	Description
Input	FAB	DNA	Default file specification string address
		DNS	Default file specification string size
		FNA	File specification string address
		FNS	File specification string size
		FOP	File-processing options (OFF bit only)
		IFI	Internal file identifier (must be zero)
		NAM	Name block address
	NAM	ESA	Expanded string area address
		ESS	Expanded string area size
		RLF	Related file NAM block address
	Related file NAM block (if any)	RSA	Resultant string area address
		RSL	Resultant string length
	Output	FAB	DEV
SDC			Spooling device characteristics
STS			Completion status code (also returned in Register 0)
STV			Status value
NAM		DID	Directory identification
		DVI	Device identification
		ESL	Expanded string length
		FID	File identification (zeroed)
		FNB	File name status bits; contains information about the parse results
		WCC	Wildcard context (zeroed to initialize the wildcard context for subsequent directory searches)

The expanded file specification string is moved to the buffer described by the expanded string area address (ESA) and size (ESS) fields of the NAM block (only if both fields are nonzero). The ESA and ESS NAM block parameters must be specified (nonzero) for wild card character processing (see Sections 7.2.2 and 7.2.3).

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

The VAX-11 completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Parse service are listed below:

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_DNF Directory not found

RMS\$_DNR Device not ready

RMS\$_WCC Invalid wild card character in NAM block

\$REMOVE

13.3 REMOVE A FILE NAME

The \$REMOVE macro instruction invokes the Remove service, which deletes a file name from a directory. (This service does not delete the file itself. The deletion is performed by the Erase service; see Section 9.4). The functions of the Remove service are performed automatically as part of an Erase service that specifies a directory.

When you remove a file name from a directory, no file can already be open for the FAB. In addition, you normally call the Parse service to set the NAM block contents before you call the Remove service.

Each removal deletes the next directory entry whose file name, type, and version number matches those specified in the expanded string length and expanded string area address fields of the NAM block.

Format

OPERATION	PARAMETERS
label: \$REMOVE	FAB=fab-address ERR=entry SUC=entry

label

A symbolic address for the \$REMOVE macro instruction; optional.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 13-3 lists the fields in both the FAB and NAM blocks that the remove service uses as input and output.

Table 13-3
Remove Fields

Usage	Control Block	Field Name	Description
Input	FAB	FOP	File-processing options (NAM bit only)
		IFI	Internal file identifier (must be zero)
		NAM	Name block address
	NAM	DID	Directory identification of directory cataloging file to be removed
		DVI	Device identification of device containing directory from which file is to be removed
		ESA	Expanded string area address specifying the name, type, and version of file to be removed
		ESL	Expanded string length
		FID	File identification; if nonzero and NAM bit is set in file-processing options field of input FAB, the first file in the directory with this file identification is removed
		FNB	File name status bits (wildcard bits only)
		RSA	Resultant string area address specifying the name, type, and version number of last file removed (required for wildcard processing)
		RSL	Resultant string length
		RSS	Resultant string area size
		WCC	Wildcard context
Output	FAB	STS	Completion status code (also returned in Register 0)
		STV	Status value
	NAM	RSL	Resultant string length
		WCC	Wildcard context

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

The resultant string is moved to the buffer described by the resultant string area address (RSA) and size (RSS) fields of the NAM block (only if both fields are nonzero).

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Remove service are listed below.

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_DNF Directory not found
RMS\$_DNR Device not ready
RMS\$_FNF File not found
RMS\$_NMF No more files found
RMS\$_PRV File protection violation
RMS\$_WLK Device write-locked

\$RENAME

13.4 RENAME A FILE

The \$RENAME macro instruction invokes the Rename service, which changes the name of a file in a directory. This service performs the equivalent of two Parse services (old and new name), a Search service for the old directory, an Enter service to insert the new file name into the new directory, and a Remove service to delete the old file name from the old directory.

When you change the name of the file in a directory, no file can already be open for the FAB, and no wild card character specifications are allowed. You can rename a file from one directory to another, but both directories must be on the same disk device.

If the Rename service is successful, the new directory entry is created and the old entry is deleted. If the service fails, the old entry remains, and the new entry, depending on when the error occurs, may or may not be created.

Format

OPERATION	PARAMETERS
label: \$RENAME	OLDFAB=fab-address ERR=entry SUC=entry NEWFAB=new-fab-address

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

label

A symbolic address for the \$RENAME macro instruction; optional.

OLDFAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB that specifies the old file name. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

NEWFAB=new-fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB that specifies the new file name. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

NOTE

If you issue this macro instruction without parameters, you must construct an additional field within your argument list to contain the address of the FAB that specifies the new file name. This additional field is placed in the argument list following the field for the success completion routine (see Section 8.1), and the argument count is set to 4.

Table 13-4 lists the fields in two FABs and two NAM blocks that the Rename service uses as input and output. In the table these blocks are called FAB#1 and NAM#1 for the old entry, and FAB#2 and NAM#2 for the new entry. For output, FAB#2 is not used, although it must be in writeable memory.

The resultant file specification string for each of the names (old and new) is placed in the buffer described by the resultant string area address (RSA) and size (RSS) fields of the separate NAM blocks (only if both fields are nonzero).

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

Table 13-4
Rename Fields

Usage	Control Block	Field Name	Description
Input	FAB#1 and FAB#2	DNA	Default file specification string address
		DNS	Default file specification string size
		FNA	File specification string address
		FNS	File specification string size
		IFI	Internal file identifier (must be zero)
		NAM	Name block address
	NAM#1 and NAM#2	ESA	Expanded string area address (must be nonzero)
		ESS	Expanded string area size (must be nonzero)
		RLF	Related file NAM block address
		RSA	Resultant string area address
		RSS	Resultant string area size
	Related file NAM blocks	RSA	Resultant string area address
		RSL	Resultant string length
	Output	FAB#1	STS
STV			Status value
NAM#1 and NAM#2		DID	Directory identification
		DVI	Device identification
		ESL	Expanded string length
		FID	File identification
		FNB	File name status bits
		RSL	Resultant string length
		WCC	Wildcard context

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Rename service are listed below.

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_DNF Directory not found
RMS\$_DNR Device not ready
RMS\$_ENT ACP enter function failed
RMS\$_FEX File already exists; not superseded
RMS\$_FNF File not found
RMS\$_IDR Invalid directory rename operation
RMS\$_PRV File protection violation
RMS\$_NMF No more files to be renamed

\$SEARCH

13.5 SEARCH FOR FILE NAME

The \$SEARCH macro instruction invokes the Search service, which scans a directory file and fills in various NAM block fields. Normally, you precede the Search service with the parse service to initialize the NAM block appropriately. The basic functions of the Search service are performed automatically as part of the Open, Create, and Erase service.

When you scan a directory file, no file can already be open for the FAB.

When called, the Search service scans the directory file specified by the directory identification field of the NAM block. It looks for an entry that matches the file name, type, and version number specified by the expanded string area address and expanded string length fields. Upon finding a match, VAX-11 RMS returns the file name, type, and version number in the buffer described by the resultant string area address and size fields, and the file identification field is filled in, thereby allowing a subsequent open by NAM block (see Section 8.2.3).

VAX-11 RMS can use wild card characters to parse through the search routine until RMS\$_NMF (No-More-Files) is reached. When the RMS\$_NMF condition is encountered, internal data structures are released.

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

Format

OPERATION	PARAMETERS
label: \$SEARCH	FAB=fab-address ERR=entry SUC=entry

label

A symbolic address for the \$SEARCH macro instruction; optional.

FAB=fab-address

Required if you use parameters in the macro instruction. This parameter defines the symbolic address of the FAB for the file. If you omit this parameter, no other parameters are permitted; you must supply the argument list within your program (see Section 8.1).

ERR=entry

The symbolic address of a user-written error completion routine; optional.

SUC=entry

The symbolic address of a user-written success completion routine; optional.

Table 13-5 lists the fields in both the FAB and NAM block that the search service uses as input and output.

The resultant file specification string is placed in the buffer described by the resultant string area address (RSA) and size (RSS) fields of the NAM block (only if both fields are nonzero). The RSA and RSS NAM block parameters must be specified (nonzero) for wild card character processing (see Sections 7.2.4 and 7.2.5).

The VAX-11 RMS completion status codes are listed in Appendix A. However, to help you anticipate any nonsevere conditions that can arise, the error or warning completion status codes for conditions that can cause a failure for the Search service are listed below.

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_DNF Directory not found
 RMS\$_DNR Device not ready
 RMS\$_FND ACP find function failed
 RMS\$_FNF File not found
 RMS\$_NMF No more files found
 RMS\$_PRV File protection violation
 RMS\$_WCC Invalid wild card context value in NAM block

FILE SPECIFICATION PROCESSING MACRO INSTRUCTIONS

Table 13-5
Search Fields

Usage	Control Block	Field Name	Description
Input	FAB	IFI	Internal file identifier (must be zero)
		NAM	Name block address
	NAM	DID	Directory identification of directory to be searched
		DVI	Device identification of device containing directory to be searched
		ESA	Expanded string area address, specifying file name, type, and version of file
		ESL	Expanded string length
		FNB	File name status bits (wildcard bits only)
		RSA	Resultant string area address, specifying name, type and version of last file found (required for wildcard processing)
		RSL	Resultant string length
		RSS	Resultant string area size
		WCC	Wildcard context
Output	FAB	STS	Completion status code (also returned in Register 0)
		STV	Status value
	NAM	FID	File identification
		RSL	Resultant string length
		WCC	Wildcard context

CHAPTER 14

RUN-TIME CONTROL BLOCK INITIALIZATION

VAX-11 RMS provides run-time equivalents of the assembly-time macro instructions that allocate and initialize control blocks. These run-time instructions are the "store" macro instructions.

The store macro instructions copy either the contents of a location or a value into data fields in the designated control block. Regardless of field size, you can access a data field with these macro instructions.

14.1 THE STORE MACRO INSTRUCTIONS

You form the name for each store macro instruction by adding `_STORE` to each assembly-time macro instruction.

For example, the run-time equivalent of the `$FAB` macro instruction is:

`$FAB_STORE`

A run-time equivalent exists for each of the following macro instructions:

- `$FAB`
- `$RAB`
- `$NAM`
- `$XABDAT`
- `$XABALL`
- `$XABKEY`
- `$XABPRO`
- `$XABFHC`
- `$XABRDT`
- `$XABSUM`

RUN-TIME CONTROL BLOCK INITIALIZATION

Format

OPERATION	PARAMETERS
label: macro-name	$\left. \begin{array}{l} \text{FAB} \\ \text{RAB} \\ \text{NAM} \\ \text{XAB} \end{array} \right\} =\text{address}$ keyword-1=value-1, ..., keyword-n=value-n

label

A user-specified symbolic address referring to the store macro instruction; optional.

macro-name

The name of the control block (FAB, RAB, NAM, XABDAT, XABALL, XABKEY, XABPRO, XABFHC, or XABRDT). The control block name is prefixed with a dollar sign (\$) and followed by `_STORE`.

address

An optional pointer to the control block; the keyword to the left of the equal sign (FAB, RAB, NAM or XAB) indicates the type of control block that you are using. The keyword XAB is used for all the different XABs.

If you specify a register name, the register must contain the address of the control block. If you specify any other type of value, the address that value represents is moved to Register 0, and that register is then used as the address of the control block. Register 0 is not preserved.

If you omit this parameter, VAX-11 RMS assumes that you have already stored the address of the control block in Register 0.

keyword-1=value-1, ..., keyword-n=value-n

A variable number of keywords that correspond to the data fields of the control block, and the values to be placed in these data fields. These values can be either keywords for options, as in the assembly-time macro instructions; or can be run-time addressing expressions. If the value is an addressing expression, the following restrictions apply:

1. For any address field -- such as the extended attribute block field (XAB) of the FAB, the file access block field (FAB) of the RAB, or the expanded string area address (ESA) and resultant string area address (RSA) fields of the NAM block -- a MOVAL instruction is generated rather than a MOVL instruction.
2. For a quadword field whose source is a register two successive registers are accessed. Therefore, the source register should not be greater than Register 11.

RUN-TIME CONTROL BLOCK INITIALIZATION

3. For any of the following fields whose source is a register, two successive registers are accessed:

- Directory identification (DID)
- File identification (FID)
- Record's file address (RFA)

Therefore, the source register should not be greater than Register 11. In addition, you cannot use the byte, word, or longword displacements for an offset, or any indexed or deferred addressing.

4. If you specify the device identification field (DVI), the source cannot be a register, since four registers would have to be accessed. In addition, you cannot use the byte, word, or longword displacements for an offset, or any indexed or deferred addressing.
5. The file protection (PRO) and group/member number (UIC) fields can be expressed in either of two ways:
 - a. Individually -- in a manner similar to the assembly-time macro instructions. For the file protection field (PRO), the values must still be the keywords R, W, E, D. For the group/member number (UIC) fields, the values must be either run-time values or constants. The radix for constants is octal.
 - b. Together -- filled in as one entity, by specifying one run-time address.

An example of a store macro instruction follows:

```
$FAB_STORE FAB=R1,ORG=SEQ,RFM=VFC,MRS=10(R2),FSZ=#30,FOP=#0,NAM=NBLK
```

In this example, Register 1 contains the address of the FAB; the file organization is sequential; the record format is variable with fixed control; and the maximum record size is to be taken from the contents of the location specified by 10(R2). In addition, the fixed size of the record is 30 bytes, the file-processing options (FOP) field is to be cleared, and the address of NBLK is to be moved into the NAM block address field of the FAB.

CHAPTER 15
CONTROL ROUTINES

VAX-11 RMS provides three control routines, as follows:

- Rundown control routine
- Default directory control routine
- Default file protection control routine

These control routines all operate synchronously; therefore, no \$WAIT macro instruction is needed.

You do not call a control routine with a macro instruction. Rather, you provide an argument list and call VAX-11 RMS at the entry point for the routine. These routines do not reference fields in the user control blocks.

15.1 HALT I/O AND CLOSE FILES

The rundown control routine closes all files opened by VAX-11 RMS for the image or process and halts I/O activity. This is not the same as closing the files with a Close service, which guarantees that all I/O will be completed (see Section 9.1). Each call made to a rundown control routine closes at least one file. Therefore, you should continue to call rundown control routines until you receive the success completion status code of RMS\$_NORMAL.

The entry point for this control routine is:

SYS\$RMSRUNDWN

There are two arguments for this control routine. The first is the address of a descriptor pointing to a 22-byte buffer to receive the device identification (16 bytes) and file identification (6 bytes) of an improperly closed output file.

The second argument is a single byte code specifying the type of rundown to be performed. This type code has the following values and meanings:

- 0 - rundown of image and indirect I/O for process permanent files
- 1 - rundown of image and process permanent files; the caller's mode must be other than user
- 2 - abort VAX-11 RMS I/O; the caller's mode must be either executive or kernel

CONTROL ROUTINES

The completion status codes are listed below.

Success:

RMS\$_NORMAL All files closed

Failure:

RMS\$_CCF An output file could not be closed successfully; user buffer identifies the file

RMS\$_IAL An output file could not be closed successfully, and the user buffer cannot be written

15.2 SET DEFAULT DIRECTORY

The default directory control routine informs you of changes in the default directory for the process. The entry point for this control routine is:

SYSS\$SETDDIR

The argument list consists of three parameters, all optional. The first is the address of the descriptor for the new default directory (or 0 if it is not to be changed). The second parameter is the address of a word to receive the length of the current default directory (or 0 if not wanted). The third is the address of the descriptor of a buffer to receive the current default directory string (or 0 if it is not wanted).

The new directory name string is checked for correct syntax.

You should restore the old default directory string to its original status unless you want the changed default directory string to last beyond the exit of your image.

The completion status codes are listed below.

Success:

RMS\$_NORMAL Operation successful

Failure:

RMS\$_DIR Directory string invalid

RMS\$_IAL Invalid argument list

15.3 SET DEFAULT FILE PROTECTION

The default file protection control routine informs you of and/or changes the default file protection for the process. The entry point for this control routine is:

SYSS\$SETDFPROT

CONTROL ROUTINES

The argument list consists of two parameters, both optional. The first is the address of a word giving the new default file protection specification (Section 6.4 describes the file protection specification), or 0 if it is not to be changed. The second parameter is the address of a word to receive the current default file protection specification, or 0 if it is not wanted.

You should restore the old default file protection specification unless you want the changed default to last beyond the exit of your image.

The completion status codes are described below.

Success:

RMS\$_NORMAL	Operation successful
--------------	----------------------

Failure:

RMS\$_IAL	Invalid argument list
-----------	-----------------------

APPENDIX A

COMPLETION STATUS CODES

This appendix lists, in alphabetical order, the completion status codes that VAX-11 RMS can return, cross-referenced to any applicable service in which they can occur. The error codes are listed in the first part of this appendix and the success codes are listed at the end.

NOTE

1. The errors that apply to the close service do not include errors that can arise due to setting of the SCF and SPL bits in the file-processing options field of the FAB.
2. The wait service has unique errors. This service can also return any status code of the awaited operation.
3. Errors associated with output operations may not necessarily be reported as the status of that particular operation because modified I/O buffers are not always written out immediately. Such errors are reported as the status of a subsequent operation, which may be an input, output or control operation.

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																											
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait	write
RMSS_ATR 0001C0CC Severe error ⊕	Read error on file header; the status value field (STV) contains an ACP error code	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	
RMSS_ATW 0001C0D4 Severe error ⊕	Write error on file header; the status value field (STV) contains an ACP error code	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	
RMSS_BES 000181C0 Warning	Invalid escape sequence entered from terminal											•																	
RMSS_BKS 0001841C Severe error	Invalid bucket size (greater than 32) in FAB		•																										
RMSS_BKZ 00018424 Severe error	Invalid bucket size (greater than 32) in the allocation XAB. The status value field (STV) contains the XAB address		•																										
RMSS_BLN 0001842C Severe error	Invalid value in block length field	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	
RMSS_BOF 00018198 Warning	File is already at beginning of the file (backspace operation)																			•		•							
RMSS_BUG 00018434 Severe error	Internal VAX-11 RMS error detected—submit an SPR		•	•					•		•					•									•				

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_BUG_DAP 00018444 Severe error	DAP protocol violation – submit an SPR	•	•	•	•	•		•		•			•		•		•	•						•		•		•
RMSS_BUG_DDI 0001843C Severe error	Invalid default directory. Internal VAX-11 RMS error; no recovery possible – submit an SPR			•				•							•	•				•								
RMSS_CCF 0001C0DC Severe error	Cannot close file; the status value field (STV) contains an error code	•																										
RMSS_CCR 00018494 Severe error	Cannot connect RAB (only one record stream permitted for sequential files or MSE not set for indexed file)		•																									
RMSS_CDA 0001C0E4 Severe error	Cannot deliver AST; the status value field (STV) contains an error code	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
RMSS_CHG 0001849C Severe error	Attempt to change a key value when that attribute not set by the key definition XAB key option flag																										•	
RMSS_CHK 000184A4 Severe error	Index file bucket check byte mismatch. The bucket has been corrupted. STV contains VBN of bucket. Submit an SPR			•						•			•				•									•		
RMSS_CHN 0001C0EC Severe error	Channel assignment failure; the status value field (STV) contains an error code		•					•	•						•	•				•	•		•					

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	ixtvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_COD 000184AC Severe error	Invalid type code in XAB. The status value field (STV) contains the XAB address	•	•		•			•							•													
RMSS_CRC 000182E2 Error	Network DAP level CRC check failed on CLOSE	•			•																							
RMSS_CRE 0001C00A Error	File create error; the status value field (STV) contains an ACP error code		•																									
RMSS_CUR 000184B4 Severe error	No current record; operation not immediately preceded by a successful get or find service			•																				•	•			
RMSS_DAC 0001C012 Error	File deaccess error during a close service; the status value field (STV) contains an ACP error code	•																										
RMSS_DAN 000184BC Severe error	Invalid data area number in key definition XAB. The status value field (STV) contains the XAB address		•																									
RMSS_DEL 00018262 Error	Record accessed by RFA record access mode has been deleted								•			•																
RMSS_DEV 000184C4 Severe error	Bad device or inappropriate device type for operation		•				•	•						•	•				•	•		•						
RMSS_DFL 0001876C Severe error	Data bucket fill size larger than bucket size specified in key definition XAB. The status value field (STV) contains the XAB address		•																									

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																											
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nxvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait	write
RMSS_DVI 000184F4 Severe error	Invalid device identification in NAM block		•				•	•							•					•	•		•						
RMSS_ENT 0001C01A Error	Error entering file name in directory; the status value field (STV) contains an ACP error code						•														•								
RMSS_ENV 00018724 Severe error	Environment error; the code necessary to support the file organization of facility was not selected at system generation		•												•														
RMSS_EOF 0001827A Error	End of file									•		•						•				•		•				•	
RMSS_ESA 000184FC Severe error	Invalid expanded string area address in NAM block		•				•	•						•	•				•	•		•							
RMSS_ESL 00018714 Severe error	Invalid expanded string length in NAM block						•												•			•							
RMSS_ESS 00018504 Severe error	Expanded string area too short		•					•						•	•					•									
RMSS_EXP 000182C2 Error	File expiration date not yet reached	•	•																										

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_EXT 0001C022 Error	File extend error; the status value field (STV) contains an ACP error code								•								•											•
RMSS_FAB 0001850C Severe error	Invalid FAB; block identifier field incorrect	•	•	•		•	•	•	•						•	•				•	•		•					
RMSS_FAC 00018514 Severe error	Operation not allowed by the value set in the file access field of the FAB				•				•	•		•					•	•							•	•		•
RMSS_FEX 00018282 Error	File already exists			•																•								
RMSS_FLG 0001851C Severe error	Invalid combination of values in key XAB FLG field; example: CHG or NUL for primary key. The status value field, (STV) contains the XAB address.			•																								
RMSS_FLK 0001828A Error	File is locked and therefore not available			•											•													
RMSS_FNA 00018524 Severe error	Invalid file specification string address in FAB			•					•						•	•				•								
RMSS_FND 0001C02A Error	Files-11 find function failed; the status value field (STV) contains an ACP error code																									•		

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	ixtvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_FNF 00018292 Error	File not found							•	•						•					•	•		•					
RMSS_FNM 0001852C Severe error	Syntax error in file name		•						•						•	•					•							
RMSS_FOP 0001853C Severe error	Invalid file processing options		•												•													
RMSS_FSZ 00018534 Severe error	Invalid fixed control area size in FAB (equal to 1 for print files)		•																									
RMSS_FTM 000187C4 Severe error	Network DAP file transfer mode does not permit operation			•						•			•					•	•							•		•
RMSS_FUL 00018544 Severe error	Device full; cannot create or extend file		•							•								•										
RMSS_IAL 0001854C Severe error	Invalid argument list	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
RMSS_IAN 00018554 Severe error	Invalid index area number in key defi- nition XAB. The status value field (STV) contains the XAB address		•																									
RMSS_IBF 00018754 Severe error	Invalid bucket format, STV contains bucket VBN – submit an SPR			•						•			•				•									•		

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																											
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait	write
RMSS_IBK 0001877C Severe error	Bucket size of lowest level of index area number (LAN) not equal to that of specified index area number (IAN field) in key definition XAB. The status value field (STV) contains the XAB address			•																									
RMSS_IDR 000182F2 Warning	Invalid directory rename operation																				•								
RMSS_IDX 0001855C Severe error	Index not initialized; internal VAX-11 RMS error. Submit an SPR																•									•			
RMSS_IFA 0001C124 Severe error	Invalid file attributes, file header corrupted; check the status value field (STV) for additional information		•	•											•														
RMSS_IFI 00018564 Severe error	Invalid internal file identifier in FAB	•	•	•		•	•	•	•						•	•				•	•		•						
RMSS_IFL 00018764 Severe error	Index bucket fill size larger than bucket size specified in key definition XAB. The status value field (STV) contains the XAB address			•																									
RMSS_IMX 0001856C Severe error	More than one XAB of the same type or non-dense XAB is present for the file. The status value field (STV) contains the XAB address	•	•			•			•						•														

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_IOP 00018574 Severe error	Invalid operation attempted: 1. block I/O when not block I/O access 2. record I/O when block I/O access 3. rewind of process permanent file 4. inappropriate device type or file organization			•				•	•	•			•	•			•	•			•		•	•	•			•
RMSS_IRC 0001857C Severe error	Invalid record encountered in file; invalid count or control byte field. The status value field (STV) contains the virtual block number for sequential and indexed files, or the relative record number for relative files. Submit an SPR			•						•			•		•		•									•		
RMSS_ISI 00018584 Severe error	Invalid internal stream identifier in RAB			•	•					•	•	•	•	•			•	•	•		•		•	•	•	•	•	•
RMSS_KBF 0001858C Severe error	Invalid key buffer address; not in access limits									•			•				•											
RMSS_KEY 00018594 Severe error	Invalid record key for random operation to a relative file. Invalid packed decimal key for an indexed file									•			•				•											
RMSS_KNM 00018774 Severe error	Invalid key name buffer address in key definition XAB. The status value field (STV) contains the XAB address		•			•									•													
RMSS_KRF 0001859C Severe error	Invalid key of reference in KRF field	•								•			•								•							

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_KSI 00018784 Severe error	Key size too large to permit two keys in index bucket, STV value is key of reference for index		•												•													
RMSS_KSZ 000185A4 Severe error	Key size not equal to 4 (relative file) or key size too large (indexed file)								•			•					•											
RMSS_LAN 000185AC Severe error	Invalid index lowest-level-bucket area number in key definition XAB. The status value field (STV) contains the XAB address		•																									
RMSS_LEX 0001878C Severe error	Attempt to extend area containing an unused extent								•																			
RMSS_LNE 000185BC Severe error	Logical name error; resulted in recursion or invalid process permanent file equivalence string		•					•						•	•					•								
RMSS_MBC 00018734 Severe error	Invalid multi-block count; must not be greater than 127	•																										
RMSS_MKD 0001C032 Error	Files-11 could not mark file for deletion; the status value field (STV) contains an ACP error code	•						•																				
RMSS_MRN 000185CC Severe error	Invalid value for maximum record number (negative) or relative key greater than maximum record number		•							•			•											•				

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	ixtvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_MRS 000185D4 Severe error	Invalid value for maximum record size			•																								
RMSS_NAM 000185DC Severe error	Invalid NAM block	•		•				•	•						•	•				•	•		•					
RMSS_NEF 000185E4 Severe error	Attempt to use the put service to a sequential file when not positioned to end of file																•											
RMSS_NET 0001874C Severe error	Network operation failed; the status value field contains DAP code	•	•	•	•	•			•		•			•				•	•					•		•		•
RMSS_NMF 000182CA Error	No more files for a search or remove operation																			•			•					
RMSS_NOD 000185F4 Severe error	Node name error			•					•							•	•				•							
RMSS_NPK 000185FC Severe error	No primary key defined in key definition XAB when creating an indexed file			•																								
RMSS_ORD 00018604 Severe error	Chained XABs not in correct (ascending) order, not dense (sequential) when required, or different types of XABs are interleaved in the same XAB sub-chain. The status value field (STV) contains the XAB address			•																								

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																											
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	INXVOL	Open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait	write
RMSS_ORG 0001860C Severe error	Unknown file organization			•												•													
RMSS_PBF 00018614 Severe error	Invalid prompt buffer address									•			•					•											
RMSS_PES 000181C8 Warning	Partial escape sequence entered from terminal; buffer overflowed before completion												•																
RMSS_PLG 0001861C Severe error	Error in file prologue; file is corrupted	•				•			•	•			•		•		•										•		
RMSS_PLV 0001872C Severe error	Prologue version unsupported														•														
RMSS_POS 00018624 Severe error	Invalid key position (greater than MRS) in key definition XAB. The status value field (STV) contains the XAB address.			•																									
RMSS_PRV 0001829A Error	Insufficient privilege or file protection violation; access denied	•		•			•	•	•						•		•			•	•		•						•
RMSS_QUO 00018634 Severe error	Error in quoted string			•					•						•	•					•								

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_RAB 0001863C Severe error	Not a valid RAB; block identifier field incorrect	•		•	•						•	•	•	•	•		•	•	•		•		•	•	•	•	•	•
RMSS_RAC 00018644 Severe error	Invalid value in record access mode field of RAB									•					•		•											
RMSS_RAT 0001864C Severe error	Record attributes invalid in FAB		•												•													
RMSS_RBF 00018654 Severe error	Invalid record address																•								•		•	
RMSS_REF 0001875C Severe error	Invalid key of reference in XAB, greater than number in file, equal to 255		•			•				•			•		•													
RMSS_RER 0001C0F4 Severe error	File read error; the status value field (STV) contains an ACP error code									•			•				•	•							•	•		
RMSS_REX 000182A2 Error	Record already exists; in a random record access mode operation to a relative file a record was found in the target record cell																•											
RMSS_RFA 0001865C Severe error	Invalid record's file address contained in RAB									•			•															

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_RFM 00018664 Severe error	Invalid record format		•	•																								
RMSS_RHB 0001866C Severe error	Invalid record header buffer									•			•				•										•	
RMSS_RLF 00018674 Severe error	Invalid related file			•					•						•	•					•							
RMSS_RLK 000182AA Error	Record locked by another process, or another stream within your process									•			•					•										
RMSS_RMV 0001C0FC Severe error	Files-11 remove function failed; the status value field (STV) contains an ACP error code																			•	•							
RMSS_RNF 000182B2 Error	Record not found									•			•															
RMSS_RNL 000181A0 Warning	Record not locked				•								•						•								•	
RMSS_ROP 0001867C Severe error	Invalid record option	•								•			•															

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_RPL 0001C104 Severe error	Error while reading prologue; the status value field (STV) contains an ACP error code	•	•	•	•				•	•		•		•		•										•		
RMSS_RRV 00018684 Severe error	Invalid RRV record encountered in indexed file, file may be corrupted			•						•		•				•										•		
RMSS_RSA 000182DA Error	Record stream active; an attempt was made to issue a record operation request in an asynchronous environment to a record stream that has a request outstanding			•	•					•	•	•	•	•		•	•	•		•			•	•	•	•		•
RMSS_RSL 0001873C Severe error	Resultant string length field of NAM block invalid																			•	•		•					
RMSS_RSS 00018694 Severe error	Resultant string area size is too small		•					•	•					•					•	•		•						
RMSS_RST 0001869C Severe error	Invalid resultant string area address in NAM block		•					•	•					•					•	•		•						
RMSS_RSZ 000186A4 Severe error	Invalid record size																									•		•
RMSS_RTB 000181A8 Warning	Record too large for user buffer													•														

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nx/vol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_SPL 0001C042 Error	Spool or submit command file option to a close service failed; the status value field (STV) contains an error code	•																										
RMSS_SQO 000186C4 Severe error	Operation not sequential									•		•					•											
RMSS_STR 000187BC Severe error	User structure (FAB/RAB) became invalid during the execution of a file or record operation	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
RMSS_SUP 000182D2 Error	Operation not supported; status value field contains DAP code	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
RMSS_SYN 000186D4 Severe error	Syntax error in file specification			•					•						•	•				•								
RMSS_SYS 0001C10C Severe error	Error in system QIO directive; the status value field (STV) contains the directive or QIO status code	•	•	•	•		•		•	•	•		•	•	•		•	•		•	•	•	•	•	•	•	•	•
RMSS_TMO 000181B0 Warning	Time-out period expired									•		•						•										
RMSS_TNS 000181B8 Warning	Terminal character not seen; Applies to terminal input only												•															
RMSS_TRE 000186DC Severe error	Index tree error: file is corrupted			•						•		•					•								•			

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_TYP 000186E4 Severe error	Error in file type		•					•							•	•				•								
RMSS_UBF 000186EC Severe error	Invalid user record area address											•						•										
RMSS_UPI 000187AC Severe error	SHR bit UPI not set when file sharing with FOP = BIO or FOP = BRO		•												•													
RMSS_USZ 000186F4 Severe error	Invalid user record area size											•						•										
RMSS_VER 000186FC Severe error	Error in version number		•					•							•	•				•								
RMSS_WBE 0001C12C Severe error	Error writing behind; the status value field (STV) contains an ACP error code	•			•			•	•	•		•				•	•			•			•	•	•	•	•	
RMSS_WCC 000182EA Error	Invalid wild card context value in NAM block														•						•							
RMSS_WER 0001C114 Severe error	File write error; the status value field (STV) contains an ACP error code	•			•			•	•	•		•				•	•			•			•	•	•	•	•	
RMSS_WLD 00018744 Severe error	Invalid wild card operation					•									•				•	•		•						

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																									
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	nextvol	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update
RMSS_WLK 000182BA Error	Device is write-locked	•	•	•	•		•	•	•	•	•	•		•		•	•		•		•		•	•	•		•
RMSS_WPL 0001C11C Severe error	Error while writing prologue; the status value field (STV) contains an ACP error code		•						•								•									•	
RMSS_XAB 0001870C Severe error	Not a valid XAB, not readable or writable, invalid code or length. The status value field (STV) contains the XAB address.	•	•		•				•					•													
RMSS_CONTROLC 00010651 Success	Operation completed under Control C; terminal I/O may have been truncated									•		•					•	•									•
RMSS_CONTROLO 00010609 Success	Operation completed under Control O; terminal output may have been truncated																•										•
RMSS_CONTROLY 00010611 Success	Operation completed under Control Y; terminal I/O may have been truncated									•		•					•	•									•
RMSS_CREATED 00010619 Success	File was created, not opened; used in conjunction with the CIF option		•																								

Status Code Hexadecimal Value Severity Level	Description	Applicable VAX-11 RMS Service																										
		close	connect	create	delete	disconnect	display	enter	erase	extend	find	flush	free	get	ioctl	open	parse	put	read	release	remove	rename	rewind	search	space	truncate	update	wait
RMSS_OK_RLK 00018021 Success	Record locked but ready anyway; locker set RLK bit in ROP field									•			•															
RMSS_OK_RNF 00018049 Success	Non-existent record accessed successfully (NXR bit set in ROP field)									•			•															
RMSS_PENDING 00018009 Success	Asynchronous operation not yet completed	•		•	•					•	•	•	•	•			•	•	•		•		•	•	•		•	
RMSS_SUC 00010001 Success	Operation successful (synonym for RMSS_NORMAL)	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
RMSS_SUPERSEDE 00010631 Success	Created file superseded an existing version		•																									

APPENDIX B

FILE/RECORD CONCEPTS AND FORMATS

VAX-11 RMS supports a variety of file organizations, record access modes, and record formats. The specific use of the file determines which file organization is best. The sections that follow outline the capabilities of each of the above items. Moreover, the Introduction to VAX-11 Record Management Services manual provides a complete description of these concepts.

B.1 FILE ORGANIZATIONS

File organization is the physical arrangement of the data in the file. You select the type of file organization you want when you create the file. Once a particular file organization is chosen, it remains fixed for the life of the file; you cannot change it. However, you can copy the file to another area, and in the process convert it to a different file organization (using the CONVERT utility).

VAX-11 RMS currently supports three file organizations:

- Sequential

In the sequential file organization, records are in physical sequence. Each record, except the first, has another record preceding it, and each record, except the last, has another record following it. The physical order in which records appear is identical to the order in which they are written. A file of sequential organization can contain records of either fixed or variable length.

- Relative

In the relative file organization, fixed-length positions, or cells, are created in the file beginning at the first record position and continue to the end of the file. There is no requirement, however, that every cell contain a record. Empty cells can be interspersed among cells that contain records. The relative file organization supports records that are either fixed or variable length.

- Indexed

In the indexed file organization, the location of records is transparent to your program; VAX-11 RMS completely controls the placement of records in an indexed file. Presence of keys in each of the records governs this placement. Records may be fixed-length or variable-length; if the records are variable-length, the maximum record length may be specified and no record can exceed the maximum length when the record is

FILE/RECORD CONCEPTS AND FORMATS

put in the file or updated. However, if a maximum length is not specified, records may be as large as the bucket size will allow. For additional information concerning the relation between bucket size and record length, see Chapter 5.

B.2 RECORD ACCESS MODES

The record access mode is the method of retrieving and storing records in a file. In contrast to file organization, which you cannot change once a file is created, you can use a different record access mode each time you process a record.

VAX-11 RMS provides three record access modes:

- Sequential

VAX-11 RMS supports sequential record access mode for all device types and file organizations.

When using the sequential record access mode, your program issues a series of requests for the next record. VAX-11 RMS interprets these requests in the context of the file organization. Thus, the organization of the file governs the order in which records are read or written; and the read or write continues, in a serial fashion, until processing of the file is completed. For sequential organization, VAX-11 RMS knows that every record after the first record is followed by another record until the end of the file (last record). For relative organization, VAX-11 RMS recognizes that empty cells can be interspersed among filled record cells and acts accordingly. On a read request, VAX-11 RMS ignores empty cells. For the indexed file, the presence of one or more indexes permits VAX-11 RMS to determine the order in which to process records in sequential access mode. Initially, your program must specify a key of reference (e.g., primary key, first alternate key, second alternate key, etc.) to VAX-11 RMS. Thereafter, VAX-11 RMS uses the index associated with that specified key to access records in the sequence represented by the entries in the index. Each successive record that VAX-11 RMS returns in response to a program read request contains a value in the specified key field that is equal to (when duplicate key values are allowed) or greater than that of the previous record returned.

- Record's File Address (RFA)

You can use the RFA record access mode with any file organization, but only for disk files and only for read operations.

The term "record's file address" means that every record in the file has a unique address. The type of file organization assigned to the file determines the format of this address.

The most important feature of RFA record access mode is that the RFA of any record remains constant while the record remains in the file. VAX-11 RMS returns the RFA to you in the RAB when the record is read or written. (The record must be written using some record access mode other than RFA, since RFA access is available for read operations only. The RFA, however, is returned in the RAB as an output from a write operation.) Your program can then save this RFA for use later

FILE/RECORD CONCEPTS AND FORMATS

during the current execution of the program, or for use at any subsequent time.

- Random by Key

VAX-11 RMS always supports random access by key for relative and indexed files. VAX-11 RMS also permits random access by relative record number for sequential disk files, but only if the records in the file are of fixed-length.

In random access by key, your program, not the file organization, determines the order in which record access occurs. Each program request for a record must include the key value (relative record number for relative files and key of reference for indexed files) of the particular record to be accessed. This program randomly identifies by means of the key value any record in the file, and VAX-11 RMS accesses that record. Your program can make successive requests for accessing records anywhere within the file.

Each of your program read requests in random access mode to an indexed file must specify a key value and the index (e.g., primary index, first alternate key index, second alternate key index, etc.) that VAX-11 RMS must search. When the VAX-11 RMS finds the key value in the specified index, it reads the record that the index entry points to and passes the record to your program. Random access can be accomplished on any key by any of the following methods:

1. Exact match of key values.
2. Approximate match of key values (e.g., record key value greater than the program-supplied key value, or record key value greater than or equal to the program-supplied key value).
3. Generic match of key values. Generic match is applicable to string data type keys only. A generic match is defined as a match on some number of leading characters in the key field. You determine the number specifying a search key which is smaller than the entire field.
4. Combination of approximate and generic match.

In contrast to read requests, which require a program-specified key value, program requests to write records randomly in an indexed file do not require the separate specification of a key value. All keys (primary and, if any, alternate key values) are in the record itself. When an indexed file is opened, VAX-11 RMS retrieves all key definitions stored in the file. Thus, VAX-11 RMS knows the location and length of each key field in a record. Before writing a record into the file, VAX-11 RMS examines the key values in the records, places the record in the file, and creates new entries in the alternate indexes. In this way, VAX-11 RMS ensures that the record can be retrieved by any of its key values.

The access mode may be switched while the file is being processed. A typical use of this feature is to perform a random by key access to locate a record. The access mode is then switched to sequential. Subsequent Get operations will return successive records by ascending key value.

FILE/RECORD CONCEPTS AND FORMATS

B.3 RECORD FORMATS

The record format is the way a record physically appears on the recording surface of the storage medium. VAX-11 RMS provides three different record formats.

- Fixed-length

The term fixed-length record format refers to file records that are all equal in size; each record occupies an equal amount of space.

- Variable-length

The term variable-length record format refers to file records that are not all the same size. VAX-11 RMS prefixes a count field to each record when it is written; this indicates to VAX-11 RMS how many bytes are in each individual record, and therefore the actual size of the record.

VAX-11 RMS uses two types of variable-length records:

- Disk files - V format

Contain a 2-byte binary count field prefixed to each record

- Tape files - D format

Contain a 4-byte decimal ASCII count field prefixed to each record

- Variable with fixed-length control (not supported for indexed files)

This type of record format is similar to V or D format variable-length records, except that it also contains a control area of fixed length. A fixed control area lets you construct variable-length records that contain an additional fixed-length piece of data that will always be present and will have a "loose" association with the other contents of the record. The VAX-11 Text Editor (see the VAX-11 Text Editor Reference Manual) uses this type of record, in which a line sequence number is associated with each line of text. This association is considered "loose" because each of the contents can be considered as separate for the purpose of processing, even though they are stored together.

Table B-1 summarizes the relationship between the VAX-11 RMS file organizations and their permitted record access modes and record formats.

B.4 FILES-11 DISK STRUCTURE

Files-11 is the term applied to the logical structure imposed on disk volumes. This structure provides the file access and allocation control mechanism for the volume. A disk volume is defined as an ordered set of blocks, with each block being an array of 512 eight-bit bytes.

FILE/RECORD CONCEPTS AND FORMATS

Table B-1
File Organization Relationships
with Record Access Modes and Record Formats

File Organization	Record Access Mode Permitted			Record Format Permitted		
	Sequential	Random by Key	Random by Record's File Address	Fixed	Variable	Variable with Fixed-Length Control
Sequential	Yes	No ¹	Yes ²	Yes	Yes	Yes
Relative	Yes	Yes ⁴	Yes	Yes	Yes ³	Yes ³
Indexed	Yes	Yes	Yes	Yes	Yes ⁵	No

¹ Random access by key (relative record number) for the sequential file organization is permitted only for the fixed-length record format on disk devices.

² Random access by RFA is permitted only on disk devices.

³ Variable-length records in the relative file organization are stored in fixed-length cells; the size of each cell is the size needed to store the largest record permitted in the file.

⁴ The key in relative file records is the relative record number.

⁵ A record in an indexed file may not cross bucket boundaries.

In terms of the volume as a whole, the blocks are numbered consecutively in the range of 0 through n-1, where n is the highest number of blocks available on the volume (this depends on the type of disk volume in use). The number assigned to each volume-relative block is the logical block number (LBN). In terms of the individual file on the volume, the blocks are numbered consecutively from 1 through the total number of blocks assigned to the file. The number assigned to each file-relative block is the virtual block number (VBN).

Figure B-1 shows the difference between the scheme of blocks considered at the LBN and VBN levels. Two files, A and B, occupy ten blocks. File A, in relation to the volume, occupies LBNs 10 through 19; but, in relation to a file, this file occupies VBNs 1 through 10. Assume that when file B was created, it was allocated in two different areas, or clusters, with each cluster five blocks in length. The first cluster occupies LBNs 300 through 304, and the second cluster is at LBNs 29 through 33. But when viewed as an individual file, file B occupies consecutive VBNs 1 through 10, just as does file A. Further assume that file B was allocated in two separate extents (this can be done either explicitly at the request of the user, or implicitly by VAX-11 RMS due to the lack of enough contiguous disk space or a default by extent size). Even though files A and B both have the same VBNs, the corresponding blocks are different since the VBNs relate to the block's placement within the individual file, not to the volume as a whole.

FILE/RECORD CONCEPTS AND FORMATS

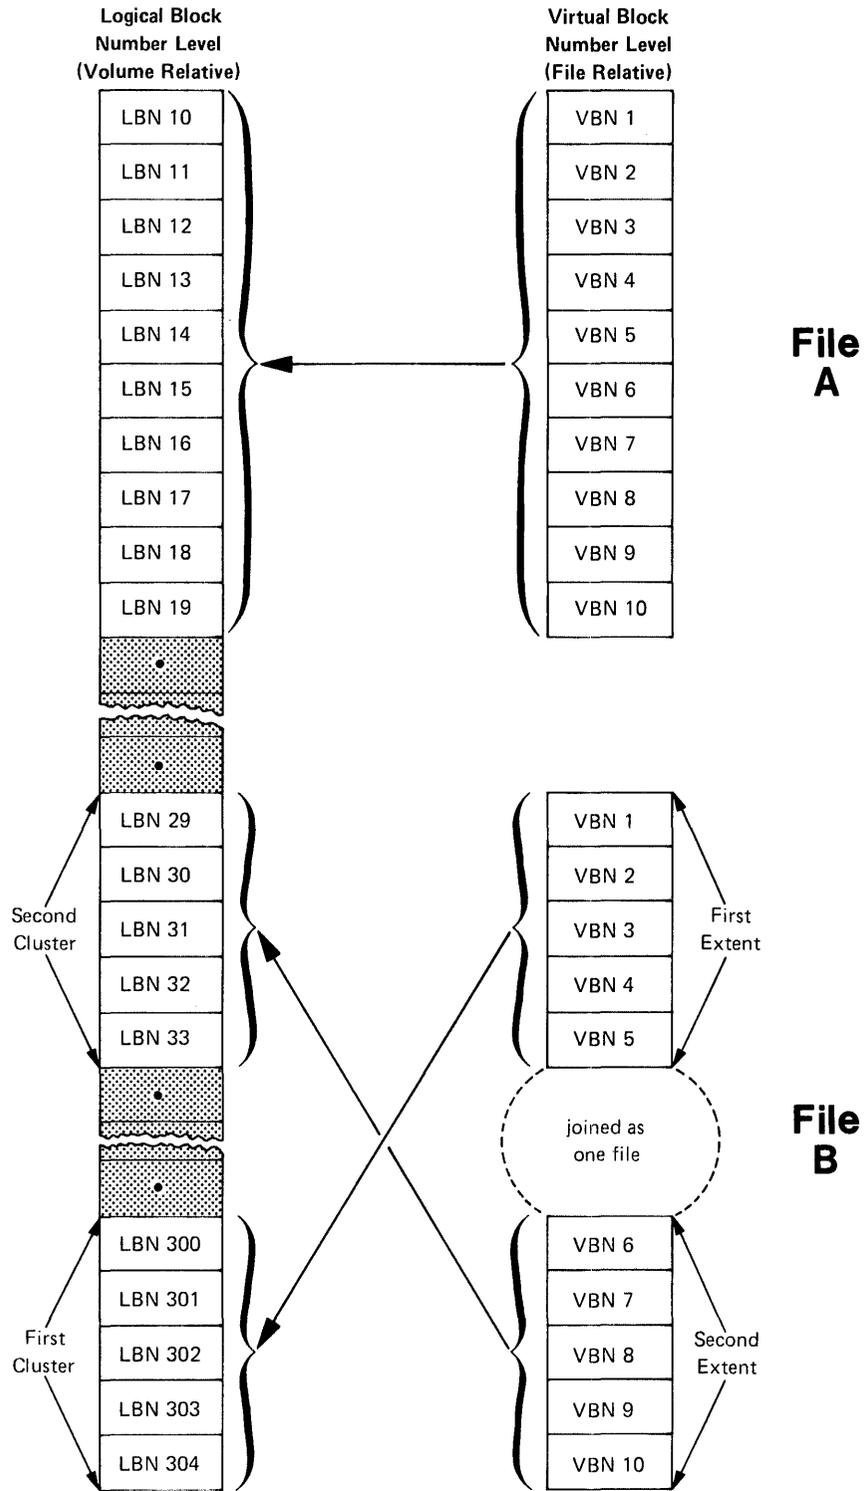


Figure B-1 Logical and Virtual Block Numbers

FILE/RECORD CONCEPTS AND FORMATS

Every Files-11 volume has an index file, which is created when the volume is initialized. This index file provides the means of identifying to VAX/VMS that the volume is a Files-11 structure, and contains the access data for all files on the volume. The index file is listed in the master file directory (MFD) as INDEXF.SYS;1 and contains the following information:

- Bootstrap block

The volume's bootstrap block is VBN 1 of the index file. Volume relative, it is LBN 0. If the volume is a system device, this block contains a program that loads the operating system into memory. If the volume is not a system device, this block contains a program that displays a message that the volume is not the system device, but rather a device that contains only user files.

- Home block

The home block identifies the volume as a Files-11 volume, establishes the specific identity of the volume, and serves as the entry point into the volume's file structure. When the volume is part of a volume set, the home block also contains the volume set name and the relative column number. The home block is VBN 2 of the index file. The LBN for the home block is the first good block (physically readable and writeable) on the volume found in the home block search sequence. The search sequence is as follows:

$$l+n * \text{delta}$$

n is in the range of 0,1,2,....

The delta is computed from the geometry of the volume such that if the volume is viewed as a three-dimensional space, the search sequence will travel down the body diagonal of the space. The dimensions included in the search delta are sectors (s), tracks (t), and cylinders (c), according to the rules in Table B-2, to handle the cases in which either one or two dimensions of the volume have a size of 1.

Table B-2
Search Delta Geometry

Geometry			Delta
s	t	c	
-	1	1	1
1	-	1	1
1	1	-	1
-	-	1	s+1
-	1	-	s+1
1	-	-	t+1
-	-	-	(t+1)*s+1

In most cases, LBN 1 will be a good block, and therefore LBN 1 will be the home block.

FILE/RECORD CONCEPTS AND FORMATS

- Back-up home block

The back-up home block is a second copy of the home block. It permits the volume to be used even if the primary home block is destroyed.

The cluster that contains the back-up home block maps into the index file at VBN $x*2+1$ through $x*3$, where x is the volume cluster factor.

- Index file bit map

The index file bit map controls (with the information contained in the home block) the allocation of file headers, and thus the number of files on the volume. The bit map contains a bit for each file header that is allowed on the volume. If the value of a bit for a given file header is 0, then a file can be created with this file header. If the value is 1, then the file header is already in use. The index file bit map starts at VBN $x*4+1$ of the index file and continues through VBN $x*4+m$, where m is the number of blocks necessary to contain the bit map, and x is the storage map cluster factor. The starting LBN for the index file bit map is recorded in the home block.

- File headers

The major portion of the index file is made up of file headers. A file header exists for each file on the volume and describes the properties of the file, such as file ownership, creation date and time, and file protection. The file header contains all the information necessary for access to the file, including the location of the file's extents.

Besides the index file, Files-11 maintains nine other files to control the volume structure. Just as with the index file, these files are created when a new volume is initialized.

The storage bit map file controls the available space on a volume, and is listed in the MFD as BITMAP.SYS;1. It contains a storage control block, which consists of summary information intended to optimize Files-11 allocation, and the bit map itself, which lists the availability of individual blocks.

The bad block file is listed in the MFD and BADBLK.SYS;1, and is simply a file containing a list of all the bad blocks on the volume.

The master file directory itself (the MFD) is listed in the MFD as 000000.DIR;1. The MFD is the root of the volume's directory structure, and lists the ten files that control the volume structure (these ten files are called the known files) plus any user files on the volume.

The core image file is listed in the MFD as CORIMG.SYS;1, and its use is operating system dependent. In general, it provides a list of the files for the operating system to use as swap areas, for example, or overlay areas.

The free space file is listed in the MFD as FREFIL.SYS;1. This file allows individual Files-11 implementations to use an alternative scheme of space allocation that is more complex than using the storage bit map file alone.

FILE/RECORD CONCEPTS AND FORMATS

The set list file is listed in the MFD as VOLSET.SYS;1. It is used only on relative volume 1 of a tightly coupled volume set. This file contains a list of the volume labels of the volumes in the volume set.

The back-up log file is listed in the MFD as BACKUP.SYS;1. It contains a history log of volume and incremental back-ups performed on this volume.

The continuation file is listed in the MFD as CONTIN.SYS;1. It is used as the extension for the file identifier when a file crosses from one volume of a loosely coupled volume set to another volume. It allows a multivolume file to be written sequentially with only one volume mounted at a time.

The pending bad block file is listed in the MFD as BADLOG.SYS;1. This file contains a list identifying suspected bad blocks on the volume that are not currently contained in the bad block file (BADBLK.SYS;1).

Each file on the volume, including the ten known files, is uniquely named by a file identifier, which is a 48-bit binary value (three words). The first word provides the file number, which locates the file on the volume. The file number is in the range of 1 through $2^{24}-1$. Once a file is deleted, its number can be reused for another file. The file number identifies the file header within the index file associated with the file. The second word is the file sequence number, which identifies the current use of a file number. This prevents any attempt to use a file identifier for a file that has already been deleted and replaced by a file with the same file number. The high byte of the third word is the relative volume number. It identifies which volume of a multivolume file contains the portion of the file that is of interest.

B.4.1 Files-11 Directories

Files-11 provides directory files to allow for accurate access to files on disk devices. A directory is a file that lists the identification and location of files owned by a particular user. Each user allowed access to a VAX/VMS system has an entry in the system authorization file defining the user identification code (UIC) and default user file directory (UFD).

Directory names can take any of three formats. Each format requires that the directory name be enclosed in either square brackets ([and]) or angle brackets (< and >). The closing bracket must match the opening bracket. The formats are as follows:

1. UIC-similar format

A UFD can be referred to in a format similar to that for a UIC: for example, [abc,xyz], where abc is a group number and xyz is the member number. This refers to a UFD of the name abcxyz.DIR;1 in the MFD. If you specify less than three characters for either abc or xyz, they are left zero-filled. Therefore, if a UFD is specified in a UIC fashion as [26,1], the directory that is searched is 026001.DIR;1 (DIR is the file type for the directory).

A UFD of this format is usually owned by a user with a corresponding UIC. This, however, is not required, since UIC and UFD ownerships are independent.

FILE/RECORD CONCEPTS AND FORMATS

2. Alphanumeric character string

A UFD can also be a 1- through 9-alphanumeric character string. This character string can be the same as your user name or account name, or any valid character string that you request or the system manager assigns you. For example, if a directory is specified as [010PAY], the directory 010PAY.DIR;1 is searched.

3. Subdirectories in addition to the character string UFD

When UFDs are referred to using the character string format, further hierarchical levels of directories can be expressed as subdirectories. A subdirectory level is expressed by adding a period (.) to the character string for the UFD, followed by the specification for the subdirectory. For example, [010PAY.DED] is the specification for the UFD named 010PAY.DIR;1 and a subdirectory of DED.DIR;1.

The maximum number of directory levels is eight: one UFD and seven subdirectories. (Combined with the master file directory, this is in effect a 9-level hierarchy.) In the directory specification [010PAY.DED.YTD], 010PAY is the UFD, DED is the first level subdirectory, and YTD is the second level subdirectory.

No maximum is placed on the number of different hierarchies of directories you can create or access.

The master file directory is created when the volume is initialized. Subdirectories and UFDs are created with the CREATE command using the DIRECTORY qualifier (see the VAX/VMS Command Language User's Guide).

The maximum number of entries that a single directory can hold ranges from 15000 to approximately 40000, depending on the length of the file specifications. In general, using several subdirectories to list a large number of files results in more efficient access than listing all files in one large directory.

The directory file itself is structured as a contiguous file with sequential organization. The records are variable-length, do not cross block boundaries, and contain no carriage control attributes.

B.5 MAGNETIC TAPE HANDLING

VAX-11 RMS support for labeled magnetic tape structure is based on the format defined by American National Standards Institute standard ANSI X3.27-1978, entitled Magnetic Tape Labels and File Structure for Information Interchange. This section describes the processing of magnetic tape files and magnetic tape labeling and file structuring format.

Magnetic tapes containing ANSI labels are coded in ASCII format, and on 9-track tape drives only.

ANSI standard X3.27-1978 allows any of the following combinations:

1. Single file on a single volume
2. Single file on more than one volume

FILE/RECORD CONCEPTS AND FORMATS

3. Multiple files on a single volume
4. Multiple files on more than one volume

Items 2 and 4 above constitute a volume set.

Magnetic tape affords sequential access only. Therefore, only one user can have access to a given volume set at any one time, and only one file in the volume set can be open for processing at a time. Access protection is performed on a volume-set basis. For volumes produced by DIGITAL systems, the owner identifier field of the volume label determines access rights (see Section B.5.1).

B.5.1 Volume Label

The volume label is always the first label on every tape volume, and serves to uniquely identify the volume and its owner. Figure B-2 presents the form of the volume label, and Table B-3 defines the contents of the fields in this label.

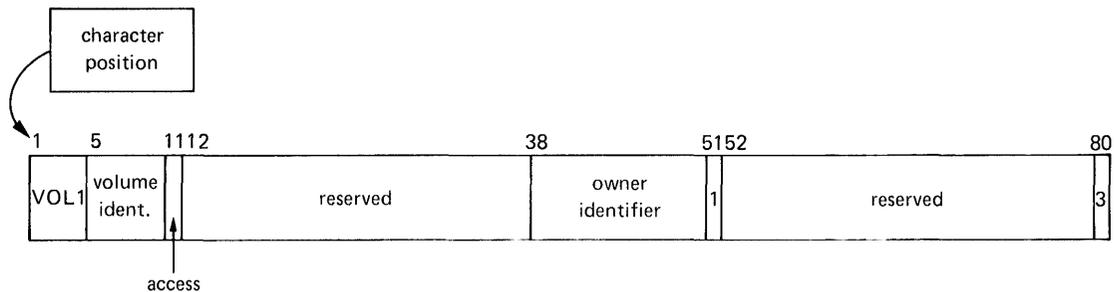


Figure B-2 Volume Label Format

FILE/RECORD CONCEPTS AND FORMATS

Table B-3
Volume Label Contents

Character Position	Field Name	Length (in bytes)	Contents
1-3	Label identifier	3	Alphabetic characters VOL
4	Label number	1	Numeric character 1
5-10	Volume identifier	6	Volume label; can be any alphanumeric or special character. This field must not be all spaces.
11	Accessibility	1	Volume protection; for the purpose of compatibility with the standards of some non-DIGITAL systems. A space (as used by DIGITAL systems) indicates no restrictions. Protects volume from being initialized.
12-37	Reserved	26	Spaces
38-50	Owner identifier	13	Volume ownership; the contents of this field are system dependent and are used for volume protection. See details following table for further amplification.
51	DIGITAL standard version	1	Numeric character 1
52-79	Reserved	28	Spaces
80	Label standard version	1	Numeric character 3

Owner identifier field

All magnetic tape volumes produced on DIGITAL systems contain the following in the first three character positions (CP 38-40) of the owner identifier field:

D%m

In the above, D% are both constant, and m represents a machine code, interpreted as follows:

- 8 - PDP-8
- A - PDP-10
- B - PDP-11
- C - VAX-11
- F - PDP-15
- K - DECSYSTEM-20

FILE/RECORD CONCEPTS AND FORMATS

If the machine code in character position (CP) 40 is the character C, the meaning of the remainder of the owner identifier field translates as follows:

1. Owner has read and write privileges:
 - CP 41-45 Group number (ASCII characters)
 - CP 46-50 Member number (ASCII characters)
2. Owner has read and write privileges; group has read privileges:
 - CP 41-45 Group number (ASCII characters)
 - CP 46 Member number high-order digit, zone encoded; therefore, a 0 in the high-order position is the character A, while a 9 is the character J
 - CP 47-50 Remaining four characters of member number (ASCII)
3. Owner has read and write privileges, world and group have read privileges:
 - CP 41 Group number high-order digit, zone encoded
 - CP 42-45 Remaining four characters of group number
 - CP 46 Member number high-order digit, zone encoded
 - CP 47-50 Remaining four characters of member number
4. Owner and group have read and write privileges:
 - CP 41-45 Group number (ASCII characters)
 - CP 46-50 Blank
5. Owner and group have read and write privileges, world has read privileges:
 - CP 41 Group number high-order digit, zone encoded; therefore, a 0 in the high-order position is the character A, while a 9 is the character J
 - CP 42-45 Remaining four characters of group number (ASCII)
 - CP 46-50 Blank
6. All categories have full privileges:
 - CP 41-50 Blank

These categories are determined when the tape is initialized using the /PROTECTION switch. Independent of what is specified in the protection code, system and owner are always granted both read and write privileges. To override this protection, either the /OWNER_UIC or /PROTECTION switch must be used at MOUNT time.

If the machine code is other than the character C, full privileges are granted unless CP11 is nonblank, in which case you must use the MOUNT command with a qualifier of /OVERRIDE=ACCESSIBILITY, to be able to initialize the tape.

FILE/RECORD CONCEPTS AND FORMATS

B.5.2 File Header Label

A file header label precedes every individual file on the tape, and serves to uniquely identify the file and describe its contents. Actually, three different file header labels precede each file; a HDR1 label, for identification, a HDR2 label, which acts as an extension to the HDR1 label and describes the characteristics of the records in the file, and a HDR3 label which contains the RMS record attributes.

Optionally, the last file header can be eliminated from files created on tape by using the /NOHDR3 switch when mounting the tape. This switch should be used when the magtape to be produced is for interchange to a system which does not tolerate HDR3 labels. The systems do not conform to the ANSI standard which requires that all labels after HDR2 be ignored on interchange tapes. Therefore, the files created on these tapes will include only HDR1 and HDR2 labels.

Figure B-3 and Table B-4 present the format and define the contents of the HDR1 label, Figure B-4 and Table B-5 present the format and define the contents of the HDR2 label, and Table B-6 describes the contents of the HDR3 label.

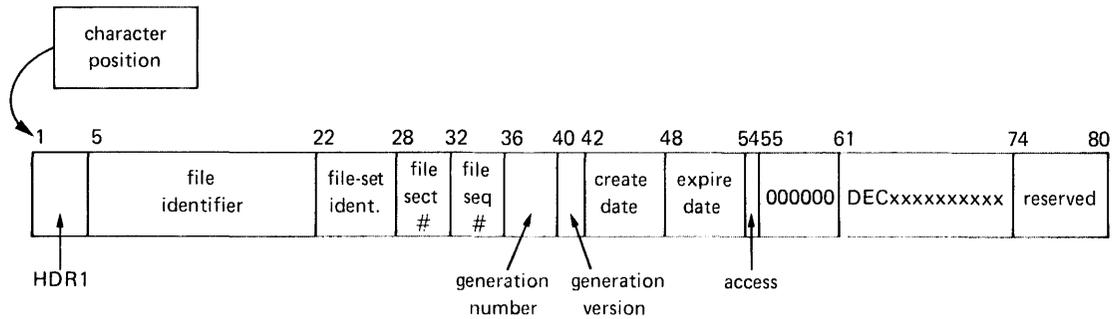


Figure B-3 HDR1 Label Format

FILE/RECORD CONCEPTS AND FORMATS

Table B-4
HDR1 Label Contents

Character Position	Field Name	Length (in bytes)	Contents
1-3	Label identifier	3	Alphabetic characters HDR to indicate a file header
4	Label number	1	Numeric character 1
5-21	File identifier	17	Any alphanumeric or special characters; see details following table for further amplification
22-27	File-set identifier	6	Same as the volume identifier of the VOL1 label of the first volume of a multi-volume set
28-31	File section number	4	Numeric characters; starts at 0001 and increments by 1 for each additional volume used by the file. This field indicates the positional order of this volume with respect to the first volume on which the file begins.
32-35	File sequence number	4	File number within the volume set for this file; consists of numeric characters, and starts at 0001. This field indicates the position of this file with respect to the first file of the set.
36-39	Generation number	4	Numeric characters; indicate the unique edition of a file. See discussion following table.
40-41	Generation version	2	Numeric characters; indicate the version number of a particular version of a file. See discussion following table.

(Continued next page)

FILE/RECORD CONCEPTS AND FORMATS

Table B-4 (Cont.)
HDR1 Label Contents

Character Position	Field Name	Length (in bytes)	Contents
42-47	Creation date	6	Julian date, in the form of yyddd (right-justified with leading space). The creation date is set to the date on which the file is created. If a creation date does not apply to this file, 00000 is used (right-justified with a leading space).
48-53	Expiration date	6	Julian date, in the form of yyddd (right-justified with a leading space). If no expiration date is specified, the value is set to the value of the creation date; therefore, the file immediately is expired.
54	Accessibility	1	File security; for the purpose of compatibility with the standards of some non-DIGITAL systems. A space (as used by DIGITAL systems) indicates no restrictions. A non-space character in this field indicates that the override switch must be used at mount time in order for the user to gain access to the file.
55-60	Block count	6	Always 000000 for the HDR1 label
61-73	System code	13	Identification code of the system that produced the file. The 3-character constant DEC appears in positions 61 through 63, followed by the name of the system. For example, DECFILE112 indicates VAX/VMS, and DECFILE11 indicates a PDP-11. The name is padded with spaces.
74-80	Reserved	7	Spaces

FILE/RECORD CONCEPTS AND FORMATS

File identifier field

The file identifier field consists of the alphabetic characters A through Z, and the numeric characters 0 through 9. ANSI standard X3.27-1978 allows special characters in this field; however, VAX/RMS translates these characters to Z.

The character preceding a period (.), or a maximum of nine characters if no period is present, constitutes the file name. The three characters following immediately after the period (or characters 10 through 12 if no period is present) constitute the file type. On output, the file name and file type are automatically separated by a period, and written to the file identifier field left-justified. The version number is generated through the generation number and generation version fields.

Generation number and generation version fields

These two fields are mapped to create the file version number, according to the following formula:

$$\text{version number} = (\text{generation number} - 1) * 100 + \text{generation version} + 1$$

For example, suppose the generation number is 11 and the generation version is 9:

$$(11 - 1) * 100 + 9 + 1$$

The formula produces a version number of 1010.

At output, the reverse is true. The present version number creates the generation number and generation version, according to the following formula and a remainder produced during the calculation.

$$\begin{array}{r} \text{generation number} = \\ \text{version number} - 1 \\ \qquad \qquad \qquad 100 \\ \qquad \qquad \qquad + 1 \end{array}$$

In the calculation, any remainder in version number - 1 is ignored for the generation number. For example, suppose the version number is 100:

$$\begin{array}{r} 100 - 1 \\ \qquad \qquad \qquad + 1 \\ \qquad \qquad \qquad 100 \end{array}$$

The formula produces a generation number of 1. The remainder of 99 is ignored in the calculation of this generation number, but becomes the generation version.

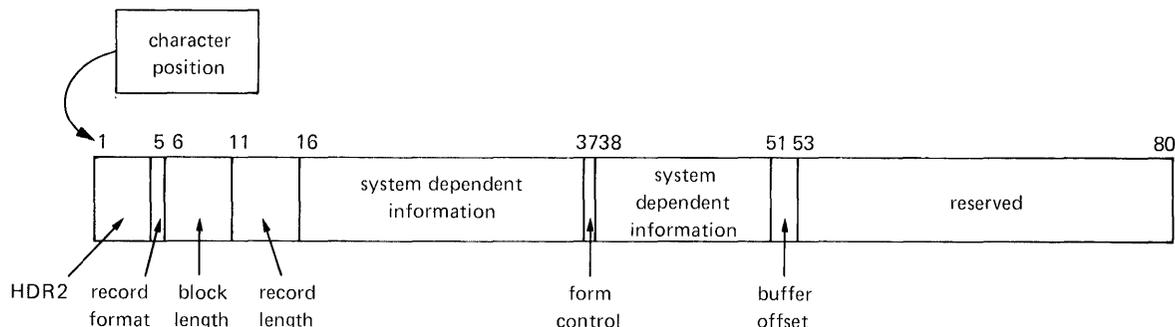


Figure B-4 HDR2 Label Format

FILE/RECORD CONCEPTS AND FORMATS

Table B-5
HDR2 Label Contents

Character Position	Field Name	Length (in bytes)	Contents
1-3	Label identifier	3	Alphabetic characters HDR to indicate a file header
4	Label number	1	Numeric character 2
5	Record format	1	<p>Character Definition</p> <p>F fixed-length D variable-length U undefined S segmented</p> <p>Undefined record format cannot be used on tapes created for interchange with non-DIGITAL systems.</p> <p>The S for segmented record formats returns as a U (undefined record format).</p>
6-10	Block length	5	Five numeric characters that specify the maximum number of characters per block.
11-15	Record length	5	Numeric characters indicating the record length for fixed-length records.
16-36	System dependent information	21	If this file was created on a VAX/VMS system, then CP's 16 through 35 contain 20 bytes of Files-11 attributes that override information in other fields of the HDR2 label; CP 36 contains a space.
37	Form control	1	<p>Defines the carriage control applied to the records in this file, as follows:</p> <p>Character Definition</p> <p>A First byte of record contains FORTRAN control characters</p>

(Continued next page)

FILE/RECORD CONCEPTS AND FORMATS

Table B-5 (Cont.)
HDR2 Label Contents

Character Position	Field Name	Length (in bytes)	Contents
			Character Definition M The record contains all form control information. space line feed/ carriage return is to be inserted between records.
38-50	System dependent information	13	If this file was created on a VAX/VMS system, then CP's 38 through 49 contain 12 bytes of Files-11 attributes that override information in other fields of the HDR2 label; CP 50 contains a space.
51-52	Buffer offset	2	The Numeric characters 00
53-80	Reserved	28	Spaces

Table B-6
HDR3 Label Contents

Character Position	Field Name	Length (in bytes)	Contents
1-3	Label Identifier	3	Alphabetic characters HDR to indicate a file header
4	Label Number	1	Numeric character 3
5-68	System-dependent information	64	If this file was created in a VAX/VMS system, then the 64 bytes contain files-11 attributes that override information in other fields of the HDR 2 label
69-80	System-dependent information	12	Spaces

B.5.3 End-of-file And End-of-volume Labels

Magnetic tape volumes contain trailer labels, which can be either of two pairs of labels, depending on whether the tape has an end-of-volume or end-of-file condition.

- End of volume

The end-of-volume label pair consists of an EOVL label and an EOVS label. These labels occur only when a file is continued

FILE/RECORD CONCEPTS AND FORMATS

onto another volume. This applies to both of the following categories of magnetic tape volumes:

- Single file, multivolume
- Multifile, multivolume

The formats of the EOVI and EOVI labels are identical to their respective HDR1 and HDR2 labels, except that the label identifier field (CP 1-3) contains EOVI and the block count field (CP 55-60) contains the number of data blocks since the last tape mark (a delimiter between labels and file data). This file data recorded since the last tape mark is known as a file section and may, in fact, be only a portion of the entire file (this occurs on a multivolume file). A file section cannot have sections of other files interspersed.

- End of file

The end-of-file label pair occurs at the end of every file recorded on a magnetic tape volume. The formats of the end-of-file labels (EOFI and EOVI) are identical to the formats of the EOVI and EOVI labels, except that the label identifier field contains EOF.

B.5.4 Arrangement Of Labels And Data

Figures B-5 through B-8 describe the organization of the different volume sets and indicate where the different labels appear. In these figures, the following legends apply:

bot = beginning of tape

*
* = tape mark
*

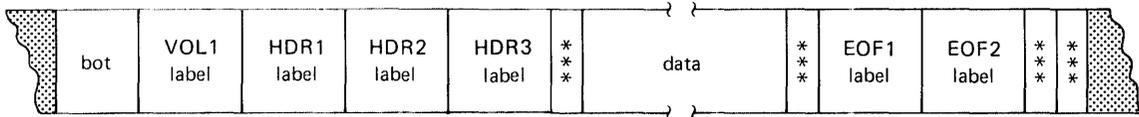


Figure B-5 Single File, Single Volume

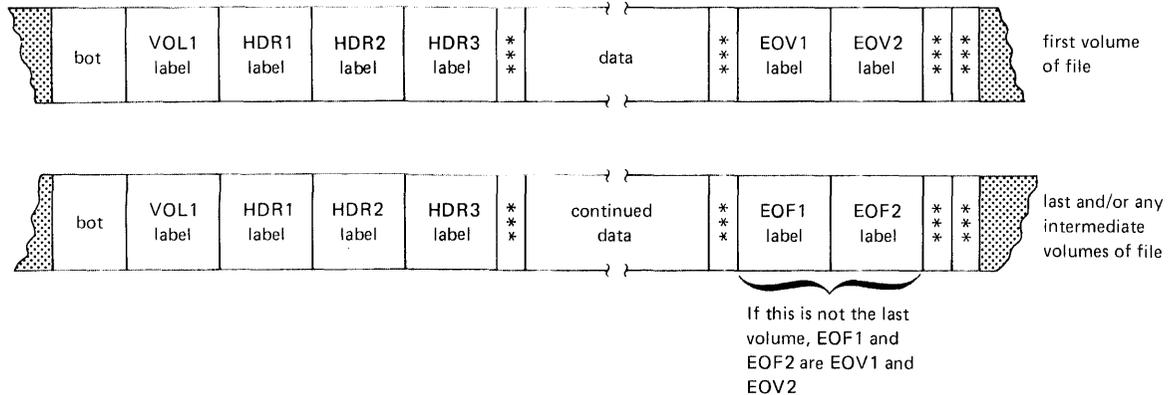


Figure B-6 Single File, Multivolume

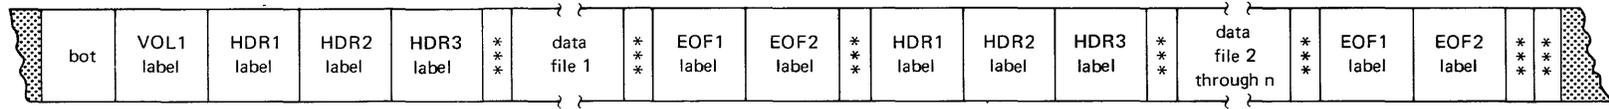
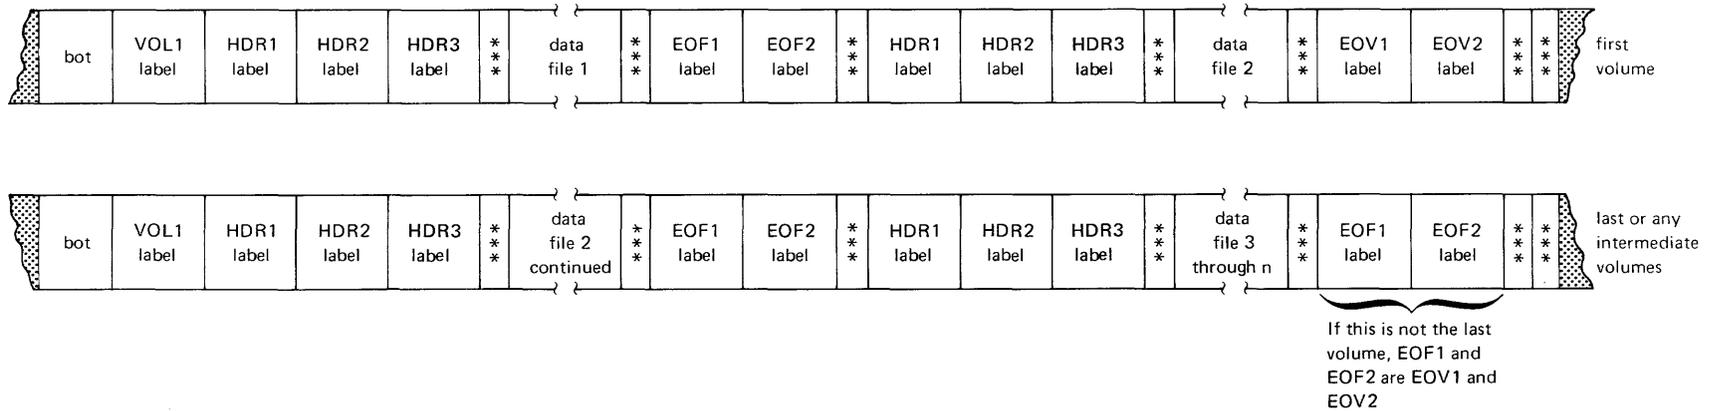


Figure B-7 Multifile, Single Volume



NOTE:

The continuation of a data file between volumes may not actually occur in data file 2; it occurs in any file which happens to be the last file on the particular volume. Data file 2 is an arbitrary choice for this figure.

Figure B-8 Multifile, Multivolume

APPENDIX C

FILE SPECIFICATION PARSING

To obtain a fully qualified file specification, VAX-11 RMS parses the primary file name string and optionally parses the default file name and related file string (if these are provided as input) as described in Section 8.2. Each of these three file name strings must have one of the following syntaxes:

- Logical-name-or-file-name
- Quoted-string-specification
- Full-file-specification

Prior to parsing a file specification, VAX-11 RMS will remove blank spaces, horizontal tabs, and null characters. If, however, such characters are placed within double-quoted strings, VAX-11 RMS will not remove them.

FILE SPECIFICATION PARSING

C.1 LOGICAL-NAME-OR-FILE-NAME SYNTAX

logical-name-or-file-name = $\left\{ \begin{array}{l} \text{logical-name} \\ \text{file-name} \end{array} \right\}$

NOTE: The logical-name takes precedence.

logical-name = $\left\{ \begin{array}{l} \text{alpha-char} \\ \text{digit} \\ \text{dollar-sign} \\ \text{underscore} \end{array} \right\} \dots$

- NOTES: 1. The logical-name is 1 through 63 alphanumeric characters, including the special characters dollar sign (\$) and underscore (_).
2. For this to be a logical name, there must be a corresponding entry in the process, group, or system logical name table.
3. If the first character of a potential logical name is an underscore, it will simply be removed by the translation process that replaces a logical name with its equivalence string. The input string, minus the leading underscore, is thus guaranteed not to be a logical name.

file-name = $\left\{ \begin{array}{l} \text{alpha-char} \\ \text{digit} \\ \text{wild-card-char} \end{array} \right\} \dots$

NOTE: The file-name is 0 through 9 alphanumeric characters, after wild card characters have been resolved. Lowercase alphabetic characters are converted to their uppercase equivalents.

THE FOLLOWING SYNTAX APPLIES TO THE
logical-name-or-file-name
TYPE OF FILE SPECIFICATION

wild-card-char = $\left\{ \begin{array}{l} * \quad (\text{asterisk}) \\ \% \quad (\text{percent sign}) \end{array} \right\}$

NOTE: The asterisk (*) wild card character may represent 0 to 9 alphanumeric characters. The percent sign (%) wild card character represents a single alphanumeric character.

FILE SPECIFICATION PARSING

C.3 FULL-FILE-SPECIFICATION SYNTAX

full-file-specification = {node-specification} {nothing}
{logical-name-or-device-name} {nothing}
{directory-specification} {nothing}
{file-name-specification} {nothing}
{file-type-specification} {nothing}
{file-version-specification} {nothing}
node-specification = see previous explanation
logical-name-or-device-name = {logical-name device-delimiter} {device-name}
NOTE: The logical-name takes precedence.
logical-name = see previous explanation
device-name = device-mnemonic {controller-name} {unit-number} {nothing} {nothing}
NOTE: For this to be a valid device name, there must be a corresponding entry in the system device data base.
device-mnemonic = uppercase-alpha uppercase-alpha NOTE: The device-mnemonic is currently limited to two characters.
controller-name = uppercase-alpha NOTE: If you omit the controller-name, the default is the character A.
unit-number = digit {digit} ... {nothing}
NOTES: 1. The unit-number is 1 through 5 digits, in the range of 0 to 65535. 2. If you omit the unit-number, the default is 0.
device-delimiter = : (single colon)
directory-specification = open-bracket directory-string close-bracket
open-bracket = { [} (left square bracket) { < } (left angle bracket)

FILE SPECIFICATION PARSING

<code>directory-string</code>	$= \left\{ \begin{array}{l} \text{group-member-form} \\ \text{directory-list} \end{array} \right\}$
<code>group-member-form</code>	$= \text{group-number group-delimiter member-number}$
<code>group-number</code>	$= \left\{ \begin{array}{l} \text{octal-number} \\ \text{wild-card-char} \end{array} \right\}$
<p>NOTE: The octal-number is 1 to 3 digits, in the range of 0 to 377 (octal).</p>	
<code>wild-card</code>	$= *$ (asterisk)
<code>group-delimiter</code>	$= ,$ (comma)
<code>member-number</code>	$= \left\{ \begin{array}{l} \text{octal-number} \\ \text{wild-card-char} \end{array} \right\}$
<p>NOTE: The octal-number is 1 to 3 digits, in the range of 0 to 377 (octal).</p>	
<code>directory-list</code>	$= \left\{ \begin{array}{l} \text{directory-name} \\ \text{wild-card-chars} \end{array} \right\} \left\{ \begin{array}{l} \text{subdirectory} \\ \text{wild-card-chars} \end{array} \right\} \dots$
<p>NOTE: You can specify a maximum of seven subdirectories</p>	
<code>directory-name</code>	$= \text{file-name}$
<p>NOTES: 1. See previous explanation of file-name. 2. If the directory-name is omitted, the current process default is used.</p>	
<code>sub-directory</code>	$= \text{directory-delimiter directory-name}$
<code>directory-delimiter</code>	$= .$ (period)
<code>wild-card-chars</code>	$= \left\{ \begin{array}{l} \text{-(minus sign)} \\ \text{...(ellipsis)} \end{array} \right\}$
<p>NOTE: The minus sign wild card character (-) represents the next highest directory level. The ellipsis wild card character (...) represents all lower directory levels.</p>	
<code>close-bracket</code>	$= \left\{ \begin{array}{l} \text{)} \text{ (right square bracket)} \\ \text{>} \text{ (right angle bracket)} \end{array} \right\}$
<p>NOTE: The close-bracket must match the open-bracket. You cannot mix square and angle brackets in the same directory-specification.</p>	
<code>file-name-specification</code>	$= \text{file-name}$
<p>See previous explanation of file-name.</p>	
<code>file-type-specification</code>	$= \text{type-delimiter file-type}$
<p>See previous explanation of wild card character.</p>	

FILE SPECIFICATION PARSING

type-delimiter	= . (period)
file-type	= { alpha-char digit wild-card-char } ...
NOTE: The file-type is 0 to 3 alphanumeric characters. Lowercase alphabetic characters are converted to their uppercase equivalents.	
wild-card-char	= { * (asterisk) % (percent sign) }
NOTE: The asterisk (*) wild card character may represent 0 to 9 alphanumeric characters. The percent sign (%) wild card character represents a single alphanumeric character.	
file-version-specification	= version-delimiter { version-number wild-card-char }
version-delimiter	= { ; (semi-colon) . (period) }
NOTE: If the version-delimiter is a period, you must include a file-type-specification.	
version-number	= { minus-sign digit nothing nothing } ...
NOTE: The version-number is 0 to 5 digits, representing a 16-bit signed value, in the range of -1 to 32767.	
wild card character	= * asterisk

APPENDIX D

DIGITAL-ONLY COMPONENT OPTIONS

There are additional file options (FOP field) which are not documented in Section 4.2.14. These options are contained in this appendix for documentation purposes only. Use of these file options by other than DIGITAL-supplied components is not supported.

DIGITAL - Only Component Options:

ESC

Escape: indicates nonstandard VAX-11 RMS processing; for DIGITAL-supplied component usage only

INP

Input: indicates that this process permanent file is the system command file named SYS\$INPUT; for DIGITAL-supplied component usage only.

KFO

Known file open: indicates a search of the known file list; for DIGITAL-supplied component usage only.

PPF

Process-permanent file: specifies that the file's internal VAX-11 RMS structures are to be allocated in the process I/O segment. The file can then be left open across images. This option applies only to DIGITAL-supplied component usage.

UFM

User file mode: indicates that the channel for the file is to be assigned in user mode. This applies only if the ESC and either the NFS or UFD options are also set. This option is provided for DIGITAL-supplied component usage only.

INDEX

\$CLOSE macro instruction, 9-1
\$CONNECT macro instruction, 11-2
\$CREATE macro instruction, 9-4
\$DELETE macro instruction, 4-12
 11-4
\$DISCONNECT macro instruction,
 11-5
\$DISPLAY macro instruction, 9-8
\$ENTER macro instruction, 13-1
\$ERASE macro instruction, 9-10
\$EXTEND macro instruction, 9-13
\$FAB macro instruction, 4-1
\$FAB_STORE macro instruction,
 4-1
\$FIND macro instruction, 4-12,
 11-8
\$FLUSH macro instruction, 11-9
\$FREE macro instruction, 11-11
\$GET macro instruction, 4-12,
 11-12
\$NAM macro instruction, 7-2
\$NXTVOL macro instruction,
 11-17
\$OPEN macro instruction, 9-14
\$PARSE macro instruction, 13-4
\$PUT macro instruction, 4-12,
 11-19
\$RAB macro instruction, 5-1
\$RAB_STORE macro instruction,
 5-1
\$READ macro instruction, 4-12,
 12-3
\$RELEASE macro instruction,
 11-22
\$REMOVE macro instruction, 13-7
\$RENAME macro instruction, 13-8
\$REWIND macro instruction, 11-24
\$SEARCH macro instruction, 13-12
\$SPACE macro instruction, 4-12,
 12-5
\$TRUNCATE macro instruction,
 4-12, 11-26
\$UPDATE macro instruction, 4-12,
 11-27
\$WAIT macro instruction, 10-8,
 11-13, 11-30
\$WRITE macro instruction, 4-12,
 12-7
\$XABALL macro instruction, 6-11
\$XABDAT macro instruction, 6-4
\$XABFHC macro instruction, 6-35
\$XABKEY macro instruction, 6-18
\$XABPRO macro instruction, 6-7
\$XABRDT macro instruction, 6-37
\$XABSUM macro instruction, 6-34
\$XABxxx STORE macro instruction,
 6-2

A

Access rights,
 delete, 6-8
 execute, 6-8
 read, 6-8
 write, 6-8
Access to process permanent files,
 8-9
AID parameter,
 area identification number
 field, 6-12
Alignment boundary type field,
 ALN parameter, 6-13
Allocation control XAB, 9-13
 \$XABALL macro instruction, 6-11
Allocation option field,
 AOP parameter, 6-14
Allocation quantity field,
 ALQ parameter, 4-4, 6-14
ALN parameter,
 alignment boundary type field,
 6-13
Angle brackets, 4-11
AOP parameter,
 allocation option field, 6-14
Area identification number field,
 AID parameter, 6-12
Argument list format,
 count, 8-2
 control block address, 8-2
 error completion routine, 8-2
 success completion routine, 8-2
Arrangement of magnetic tape
 labels, B-20
ASY record-processing option bit,
 5-13
Asynchronous operations, 8-8,
 10-1, 10-7, 10-8, 11-9,
 11-17, 11-30
Asynchronous record-processing
 option, 5-13
Automatic disk file extension,
 4-7
Automatic record locking, 10-10

B

Backup home block, B-7
Backup log file, B-9
Bad block file, B-8
BIO,
 file access option bit, 4-12
 record-processing option bit,
 5-13
BKS parameter,
 bucket size field, 4-5

INDEX

- BKT parameter,
 - bucket code field, 5-4
 - BKZ parameter,
 - bucket size field, 6-15
 - BLK bit, 4-22
 - Block, B-5
 - Block boundaries, 11-12
 - Block I/O, 4-23, 5-4, 5-8, 5-20, 12-1
 - Block I/O record-processing option, 5-14
 - Block identifier field, 4-27, 5-21
 - Block length field, 4-27, 5-20
 - Block size field,
 - BLS parameter, 4-7
 - BLS parameter,
 - block size field, 4-7
 - Bootstrap block, B-7
 - BRO file access option bit, 4-12
 - Bucket code field,
 - BKT parameter, 5-4
 - Bucket size field,
 - BKS parameter, 4-5
 - BKZ parameter, 6-16
 - Bucket size formulas, 4-6
- C**
- Cancel control 0 record-processing option, 5-17
 - CBT,
 - allocation option bit, 6-14
 - file-processing option bit, 4-15
 - CCO record-processing option bit, 5-17
 - Chained XABs, 6-3
 - CIF file-processing option bit, 4-16
 - Close service,
 - \$CLOSE macro instruction, 9-1
 - Close all files, 15-1
 - Cluster, B-5
 - Completion routine conventions, 8-3
 - Completion status code field, 4-27, 5-21, 8-8
 - Completion status codes, 8-1, 8-8, A-1
 - Connect service,
 - \$CONNECT macro instruction, 11-2
 - Contiguous file-processing option, 4-15
 - Contiguous best try file-processing option, 4-15
 - Continuation file, B-9
 - Control block,
 - access, 1-2
 - alignment, 3-1
 - Control block, (Cont.)
 - allocation, 1-1
 - initialization, 1-1
 - use, 8-1, 8-7
 - Control routines, 15-1
 - Convert record-processing option, 5-17
 - Core image file, B-8
 - CR bit, 4-22
 - Create by NAM block, 8-7
 - Create if file-processing option, 4-16
 - Create service,
 - \$CREATE macro instruction, 9-4
 - Creation data and time, 6-6
 - CTG,
 - allocation option bit, 6-15
 - file-processing option bit, 4-15
 - CTX parameter,
 - user context field, 4-8, 5-5
 - Current context of a stream, 11-24
 - Current position file-processing option, 4-17
 - Current Record,
 - contents, 10-3
 - context, 10-1, 10-3
 - CVT record-processing option bit, 5-17
- D**
- D format variable-length records, B-4
 - DAN parameter,
 - data buckets area number field, 6-20
 - Data buckets area number field,
 - DAN parameter, 6-20
 - Data buckets fill size field,
 - DFL parameter, 6-21
 - Date and time extended attribute block fields, 6-4
 - Data and time XAB,
 - \$XABDAT macro instruction, 6-4
 - Declaring manual record locking, 10-9, 10-11
 - Default directory control routine, 15-1
 - Default extension quantity field,
 - DEQ parameter, 4-9, 6-16
 - Default file protection control routine, 15-1, 15-3
 - Default file specification string address field,
 - DNA parameter, 4-10
 - Default file specification string size field,
 - DNS parameter, 4-11

INDEX

- Deferred write file-processing option, 4-15
 - Definition of terms, 1-2
 - DEL,
 - file access option bit, 4-12
 - file-sharing bit, 4-25
 - Delete access rights, 6-9
 - Delete file-processing option, 4-16
 - Delete service,
 - \$DELETE macro instruction, 11-3
 - Deleting a file name, 13-7
 - DEQ parameter,
 - default extension quantity field, 4-9, 6-16
 - Device characteristics field, 4-27
 - Device identification, 8-7
 - DFL parameter,
 - data buckets fill size field, 6-21
 - DFW file-processing option bit, 4-15
 - Directory,
 - entry removal, 9-10
 - file scan, 13-12
 - identification, 8-7
 - specification, 8-6, 8-7
 - Disconnect service,
 - \$DISCONNECT macro instruction, 11-5, 11-6
 - Disk volume, B-5
 - Display service,
 - \$DISPLAY macro instruction, 9-8
 - DLT file-processing option bit, 4-16
 - DNA parameter,
 - default file specification string address field, 4-11
 - DNM parameter, 4-11
 - DNS parameter,
 - default file specification string size field, 4-10
 - DTP parameter,
 - key data type field, 6-22
 - Dynamic access, 10-1
- E**
- EDT parameter,
 - expiration date and time field, 6-5
 - End of file, 11-16
 - End of file labels, B-20
 - End-of-file record-processing option, 5-14
 - End of volume labels, B-19
 - Enter service,
 - \$ENTER macro instruction, 13-1
 - EOF record-processing option bit, 5-14
 - EOF1 label, B-20
 - EOF2 label, B-20
 - EOV1 label, B-20
 - EOV2 label, B-20
 - Erase Service,
 - \$ERASE macro instruction, 9-10
 - Error status codes, 8-8
 - ESA parameter,
 - expanded string area address field, 7-3
 - ESC file-processing option bit, D-1
 - ESS parameter,
 - expanded string area size field, 7-4
 - Establishing a record stream, 11-2
 - Execute access rights, 6-8
 - Expanded string area address field,
 - ESA parameter, 7-3
 - Expanded string area size field,
 - ESS parameter, 7-4
 - Expiration date and time field,
 - EDT parameter, 6-5
 - Explicit assembly time initialization, 8-6
 - Extend service,
 - \$EXTEND macro instruction, 9-13
 - Extended attribute block chain, 6-1
 - Extended attribute block pointer field,
 - XAB parameter, 4-26
 - Extended Attribute Blocks, 4-26
- F**
- FAB,
 - allocation, 4-1, 4-3
 - fields, 4-1
 - FAB parameter,
 - file access block address field, 5-5
 - FAB parameters,
 - ALQ, 4-4
 - BKS, 4-5
 - BLS, 4-8
 - CTX, 4-9
 - DEQ, 4-9
 - DNA, 4-10
 - DNM, 4-11
 - DNS, 4-11
 - FAC, 4-12
 - FNA, 4-13
 - FNM, 4-14
 - FNS, 4-14

INDEX

- FAB parameters, (Cont.)
 - FOP, 4-14
 - FSZ, 4-18
 - MRN, 4-19
 - MRS, 4-20
 - NAM, 4-20
 - ORG, 4-21
 - RAT, 4-21
 - RFM, 4-23
 - RTV, 4-24
 - SHR, 4-25
 - XAB, 4-26
- FAC parameter,
 - file access field, 4-12
- File access bit offset, 4-13
- File access block field,
 - FAB parameter, 5-5
- File access block,
 - FAB, 4-1, 9-1
- File access field,
 - FAC parameter, 4-12
- File access mask value, 4-13
- File access option bits, 4-12
- File access privileges, 6-8
- File attribute information, 9-8
- File extension, 9-13
- File header characteristics XAB,
 - \$XABFHC macro instruction, 6-35
- File header labels, B-14
- File headers, B-8
- File identification, 8-8
- File identifier, B-9
- File name,
 - change, 13-9
 - deletion, 13-7
 - insertion, 13-1
 - rename service, 13-9
- File name status bits, 7-8
- File number, B-9
- File organization, 1-1, B-1
- File organization field,
 - ORG parameter, 4-21
- File positioning, 12-5
- File-processing macro
 - instructions, 9-1
- File-processing option bits,
 - CBT, 4-15
 - CIF, 4-16
 - CTG, 4-15
 - DFW, 4-15
 - DLT, 4-16
 - ESC, D-1
 - INP, D-1
 - KFO, D-1
 - MXV, 4-16
 - NAM, 4-16
 - NEF, 4-17
 - NFS, 4-17
 - OFP, 4-16
- File-processing option bits (Cont.)
 - POS, 4-17
 - PPF, D-1
 - RCK, 4-15
 - RWC, 4-17
 - RWO, 4-17
 - SCF, 4-16
 - SPL, 4-17
 - SQO, 4-15
 - SUP, 4-16
 - TEF, 4-15
 - TMD, 4-17
 - TMP, 4-17
 - UFM, D-1
 - UFO, 4-18
 - WCK, 4-16
- File-processing options filed,
 - FOP parameter, 4-14
- File protection field,
 - PRO parameter, 6-8
- File protection XAB,
 - \$XABPRO macro instruction, 6-1, 6-7
- File sequence number, B-9
- File sharing, 10-8
- File-sharing field,
 - SHR parameter, 4-24, 11-2
- File specification,
 - components, 4-10
 - default application, 8-4
 - parsing, 7-4, 13-4, C-1
- File specification processing
 - macro instructions, 13-1
- File specification string address
 - field,
 - FNA parameter, 4-13
- File specification string size
 - field,
 - FNS parameter, 4-14
- Files-11 directories, B-9
- Find service,
 - \$FIND macro instruction, 11-7
- Fixed control area, B-4
- Fixed control area size, 5-13
- Fixed control area size field,
 - FSZ parameter, 4-18
- Fixed-length record format, 4-23, B-4
- FLG parameter,
 - key options flag field, 6-24
- Flush service,
 - \$FLUSH macro instruction, 11-9
- FNA parameter,
 - file specification string
 - address field, 4-13
- FNM parameter, 4-14
- FNS parameter,
 - file specification string size
 - field, 4-14

INDEX

FOP parameter,
 file-processing options field,
 4-15
 FORTRAN carriage control, 4-22
 Free service,
 \$FREE macro instruction, 11-11
 Free space file, B-8
 FSZ parameter,
 fixed control area size field,
 4-18
 FTN bit, 4-22
 Fully qualified file specifica-
 tion, 8-7

G

GET,
 file access option bit, 4-12
 file-sharing bit, 4-25
 Get service,
 \$GET macro instruction, 11-12
 Group and member number field,
 UIC parameter, 6-10
 Group user class, 6-8

H

HDR1 label, B-14
 HDR2 label, B-14
 Home block, B-7

I

IAN/6
 IAN parameter,
 index buckets area number
 field, 6-26
 IFL parameter,
 index buckets fill size field,
 6-27
 Implicit assembly time
 initialization, 8-6
 Independent record stream, 10-5
 Index buckets area number field,
 IAN parameter, 6-26
 Index buckets fill size field,
 IFL parameter, 6-27
 Index file, B-7
 Index file bit map, B-8
 Indexed file, 6-15, 6-19
 INP file-processing option bit,
 D-1
 Internal stream identifier field,
 ISI, 5-20
 Internal file identifier field,
 IFI, 4-28

K

KBF parameter,
 key buffer address field, 5-5
 KFO file-processing option bit,
 D-1
 Key buffer address field,
 KBF parameter, 5-6
 Key data type field,
 DTP parameter, 6-24
 Key definition XAB,
 \$XABKEY macro instruction,
 6-21
 Key definition XAB parameters,
 DAN, 6-20
 DFL, 6-21
 DTP, 6-22
 FLG, 6-24
 IAN, 6-26
 KNM, 6-28
 LAN, 6-29
 NUL, 6-29
 POS, 6-30
 REF, 6-31
 SIZ, 6-32
 Key name address field,
 KNM parameter, 6-28
 Key position field,
 POS parameter, 6-30
 Key of reference field,
 KRF parameter, 5-7
 Key options flag field,
 FLG parameter, 6-24
 KEY record access mode bit, 5-12
 Key size field,
 KSZ parameter, 5-7
 SIZ parameter, 6-32
 Keys,
 alternate, 6-20, 6-25
 primary, 6-20, 6-26
 segmented, 6-32
 simple, 6-32
 size, 6-32
 Known file open file-processing
 option, D-1
 KNM parameter,
 key name address field, 6-28
 KRF parameter,
 key of reference field, 5-7
 KSZ parameter,
 key size field, 5-7

L

LAN parameter,
 lowest level of index area
 number field, 6-29
 LBN,
 logical block number, B-5

INDEX

LOC parameter,
 location field, 6-16
 LOC record-processing option
 bit, 5-15
 Locate mode, 5-15, 10-2, 11-13,
 11-9, 11-27
 Locate mode record-processing
 option, 5-15
 Location field,
 LOC parameter, 6-17
 Logical block numbers, B-5
 Logical names, 7-3, 8-9
 Lowest level of index area
 number field,
 LAN parameter, 6-29

M

Macro instructions,
 general format, 8-1
 Magnetic tape, B-10
 Magnetic tape interchange, 4-8
 Magnetic tape labels, B-19
 Manual unlock record-processing
 option, 5-16
 Manual record locking, 10-11
 declaration of, 10-11
 Master file directory,
 MFD, B-7, B-8
 Maximize version file-processing
 option, 4-16
 Maximum record number field,
 MRN parameter, 4-19
 Maximum record size field,
 MRS parameter, 4-19
 Maximum record sizes,
 fixed-length records, 4-19
 variable-length records, 4-19
 variable with fixed control
 records, 4-20
 MBC parameter,
 multiblock count field, 5-8
 MBF parameter,
 multibuffer count field, 5-9
 MFD,
 master file directory, B-7, B-8
 Modifying record contents, 11-27
 Move mode, 10-2, 11-17
 MRN parameter,
 maximum record number field,
 4-19
 MRS parameter,
 maximum record size field, 4-20
 MSE file-sharing bit, 4-25
 Multiblock count field,
 MBF parameter, 5-9
 Multiple record streams, 10-8
 Multistream access, 4-25
 MXV file-processing option bit, 4-16

N

NAM block,
 allocation, 7-2
 create by, 8-7
 fields, 7-2
 open by, 8-4, 8-7
 NAM block input file-processing
 option, 4-16
 NAM block parameters,
 ESA, 7-3
 ESS, 7-4
 RLF, 7-4
 RSA, 7-5
 RSS, 7-5
 NAM file-processing option bit,
 4-16
 Name block,
 NAM block, 7-1
 Name block address field,
 NAM parameter, 4-20
 NEF file-processing option bit,
 4-17
 Next block pointer, 12-3
 Next Record, 10-4
 contents, 10-4
 Next volume service,
 \$NXTVOL macro instruction,
 11-17
 Next XAB address field,
 NXT parameter, 6-3, 6-4
 NFS file-processing option bit,
 4-17
 NIL file-sharing bit, 4-25
 NLK record-processing option
 bit, 5-15
 No lock record-processing option,
 5-15
 Nonexistent record-processing
 option, 5-16
 Noninitializable FAB fields, 4-27
 Noninitializable key fields, 6-33
 Noninitializable NAM block
 fields, 7-6
 Noninitializable RAB fields, 5-21
 Nonfile-structure file-processing
 option, 4-17
 Nonfile-structured operations,
 12-10
 Not end of file-processing
 option, 4-17
 NXR record-processing option bit,
 5-16
 NXT parameter,
 next XAB address field, 6-3,
 6-5
 Null key value field,
 NUL parameter, 6-29
 NUL parameter field,
 null key value, 6-29

INDEX

O

OFP file-processing option bit,
 4-16
 Open by NAM block, 8-4, 8-7,
 9-14
 Open service,
 \$OPEN macro instruction, 9-14
 Order of chained XABs, 6-3
 ORG parameter,
 file organization field, 4-21
 Output file parse file-processing
 option, 4-16
 Owner user class, 6-8

P

Parameter delimiters, 8-2
 Parse service, 13-6, 13-12
 \$PARSE macro instruction, 13-4
 Parsing a file specification,
 13-5, C-1
 Path to a file, 4-13, 8-3
 PBF parameter,
 prompt buffer address field,
 5-10
 PMT record-processing option
 bit, 5-17
 POS file-processing option bit, 4-16
 POS parameter,
 key position field, 6-30
 Positioning a file, 12-5
 PPE file-processing option bit,
 4-16
 Primary,
 index, 6-26
 key, 6-21
 PRN bit, 4-22
 PRO parameter,
 file protection field, 6-8
 Process permanent file-processing
 option, 4-16
 Process permanent files, 8-9
 Program section \$RMSNAM, 4-11
 Prompt buffer address field,
 PBF parameter, 5-10
 Prompt buffer size field,
 PSZ parameter, 5-11
 Prompt record-processing
 option, 5-17
 PSZ parameter,
 prompt buffer size field, 5-11
 PTA record-processing option
 bit, 5-17
 Purge type-ahead record-proces-
 sing option, 5-17
 PUT,
 file access option bit, 4-12
 file-sharing bit, 4-25

Put service,
 \$PUT macro instruction, 11-9

R

RAB,
 allocation, 5-3
 fields, 5-1
 RAB parameters,
 BKT, 5-4
 CTX, 5-5
 FAB, 5-5
 KBF, 5-6
 KRF, 5-6
 KSZ, 5-7
 MBC, 5-8
 MBF, 5-9
 PBF, 5-10
 PSZ, 5-11
 RAC, 5-11
 RBF, 5-12
 RHB, 5-13
 ROP, 5-13
 RSZ, 5-18
 TMO, 5-17
 UBF, 5-20
 USZ, 5-20
 RAC parameter,
 record access mode field, 5-11
 RAH record-processing option bit,
 5-15
 Random access by key value, 5-6,
 5-11, 10-1
 Random access by record's file
 address, 5-12, 10-2, 11-22
 Random access by RFA record
 access mode, 11-20
 Random starting point, 11-7
 RAT parameter,
 record attributes field, 4-21
 RBF parameter,
 record address field, 5-12
 RCK file-processing option bits,
 4-15
 Read access rights, 6-8
 Read no echo record-processing
 option, 5-17
 Read no filter record-processing
 option, 5-17
 Read of locked records allowed,
 5-16
 Read service,
 \$READ macro instruction,
 12-3
 Read-ahead record-processing
 option, 5-15
 Read-check file-processing
 option, 4-16
 Record access, 10-1

INDEX

- Record access block,
 - RAB, 5-1, 11-1
- Record access mode,
 - random by key (relative record number), 5-6, 5-12, 10-1, B-2
- Record access mode
 - random by record's file address, 5-12, 10-2, 11-20, B-2
 - sequential, 5-12, 10-1, B-2
- Record access mode bits, 5-12
- Record access mode field,
 - RAC parameter, 5-11
 - specification, 10-1
- Record address field,
 - RBF parameter, 5-12
- Record attributes field,
 - RAT parameter, 4-21
- Record cell, B-1
- Record contents, 11-25
- Record control information, 4-21
- Record format,
 - fixed-length, B-4
 - variable with fixed-length control, B-4
 - variable-length, B-4
- Record format field,
 - RFM parameter, 4-23
- Record header field,
 - RHB parameter, 5-13
- Record-processing macro instructions, 11-1
- Record-processing options bits,
 - ASY, 5-13
 - BIO, 4-14
 - CCO, 5-17
 - CVT, 5-17
 - EOF, 5-14
 - KGE, 5-14
 - KGT, 5-14
 - LIM, 5-15
 - LOA, 5-14
 - LOC, 5-15
 - NLK, 5-16
 - NXR, 5-16
 - PMT, 5-17
 - PTA, 5-17
 - RAH, 5-15
 - RLK, 5-16
 - RNE, 5-17
 - RNF, 5-17
 - TMO, 5-17
 - TPT, 5-16
 - UIF, 5-16
 - ULK, 5-16
 - WBH, 5-15
- Record-processing options field,
 - 10-2, 10-11,
 - ROP parameter, 5-13
 - use with record locking, 10-10
- Record,
 - locking, 10-1, 10-9, 10-10, 11-7
 - removal, 11-3
 - retrieval, 11-12
- Record (cont.)
 - skipping, 11-7
 - stream, 5-1, 11-2
 - transfer mode, 5-18
 - types of, 10-10
 - unlocking, 11-11, 11-22
- Record size field,
 - RSZ parameter, 5-18
- Record streams, 10-1, 10-5
- Record's file address, 5-20
- Related file NAM block address field,
 - RLF parameter, 7-4
- Relative file organization, B-1
- Relative record number, 4-19
- Relative volume number, B-9
- Relative volume number field,
 - VOL parameter, 6-18
- Release service,
 - \$RELEASE macro instruction, 11-22
- Remove service,
 - \$REMOVE macro instruction, 13-7
- Removing records, 11-3
- Rename service,
 - \$RENAME macro instruction, 13-9
- Resultant file specification,
 - string, 7-5
- Resultant string area address field,
 - RSA parameter, 7-5
- Resultant string area size field,
 - RSS parameter, 7-5
- Retrieval window size field,
 - RTV parameter, 4-24
- Revision date and time field,
 - 6-6
- Revision date and time XAB,
 - \$XABRDT macro instruction, 6-38
- Revision number, 6-5, 6-38
- Rewind on close file-processing option, 4-16
- Rewind on open file-processing option, 4-17
- Rewind service,
 - \$REWIND macro instruction, 11-22
- RFA record access mode bit, 5-12
- RFM parameter,
 - record format field, 4-23
- RHB parameter,
 - record header field, 5-13

INDEX

- RLF parameter,
 - related file NAM block address field, 7-4
 - RLK record-processing option bit, 5-16
 - RNE record-processing option bit, 5-17
 - RNF record-processing option bit, 5-17
 - ROP parameter,
 - record-processing options field, 5-13
 - RSA parameter,
 - Resultant string area address field, 7-5
 - RSS parameter,
 - Resultant string area size field, 7-5
 - RSZ parameter,
 - record size field, 5-18
 - RTV parameter,
 - retrieval window size field, 4-24
 - Run-time,
 - control block initialization, 14-1
 - initialization, 8-6
 - processing interface, 8-1
 - Rundown control routine, 15-1
 - RWC file-processing option bit, 4-16
 - RWO file-processing option bit, 4-16
- S**
- SCF file-processing option bit, 4-16
 - Search service,
 - \$SEARCH macro instruction, 13-12
 - Segmented keys, 6-30
 - SEQ record access mode bit, 5-11
 - Sequential file organization, B-2
 - Sequential only file-processing option, 4-15
 - Sequential record access mode, 5-12, 10-1
 - Set list file, B-9
 - Shared sequential files, 4-25, 6-15
 - SHR parameter,
 - file-sharing field, 4-25
 - Simple keys, 6-32
 - Single record stream, 10-5
 - SIZ parameter,
 - key size field, 6-32
 - Skipping records, 11-7
 - Space service,
 - \$SPACE macro instruction, 12-5
 - SPL file-processing option bit, 4-17
 - Spool file-processing option, 4-17
 - Spool device characteristics field, 4-27
 - SQO file-processing option bit, 4-15
 - Statement conventions, 2-1
 - Status value field, 4-27, 5-21, A-1
 - Store macro instructions,
 - addressing expression restrictions, 14-2
 - formation, 14-1
 - Storage bit map file, B-8
 - Subdirectory, B-10
 - Submit command file-processing option, 4-16
 - Summary XAB parameter,
 - NXT, 6-34
 - SUP file-processing option bit, 4-16
 - Supersede file-processing option, 4-16
 - Synchronous operations, 10-1, 10-7
 - SY\$ERROR, 8-9
 - SY\$INPUT, 8-9
 - SY\$OUTPUT, 8-9
 - SY\$RMSRUNDWN, 15-1
 - SY\$SETDDIR, 15-2
 - SY\$SETDFPROT, 15-2
 - System service exceptions, 8-8
 - System user class, 6-8
- T**
- TEF file-processing option bit, 4-15
 - Temporary file-processing option, 4-17
 - Temporary marked for delete file processing option, 4-17
 - Terminating a record stream, 11-5
 - Time-out period field,
 - TMO parameter, 5-19
 - Time-out record-processing option, 5-17
 - TMD file-processing option bit, 4-17
 - TMO parameter,
 - time-out period field, 5-19
 - TMO record-processing option bit, 5-16
 - TMP file-processing option bit, 4-17

INDEX

TPT record-processing option bit, 5-17
 Translation of logical names, 7-3, 8-9
 TRN file access option bit, 4-12
 Truncate at end of file-processing option, 4-15
 Truncate put record-processing option, 5-16
 Truncate service,
 \$TRUNCATE macro instruction, 11-26
 Types of record locking, 10-10

U

UBF parameter,
 user record area address field, 5-20
 UFD,
 user file directory, B-9
 UFM file-processing option bit, D-1
 UFO file-processing option bit, 4-18
 UIC,
 user identification code, B-9
 UIC parameter,
 group and member number fields, 6-8
 UIF record-processing option, bit, 5-16
 ULK record-processing option bit, 5-16
 Undefined record format, 4-24
 Unlocking records, 11-11, 11-22
 UPD,
 file access option bit, 4-12
 file-sharing bit, 4-25
 Update service,
 \$UPDATE macro instruction, 11-27
 UPI file-sharing bit, 4-25
 User classes,
 group, 6-8
 owner, 6-8
 system, 6-8
 world, 6-8
 User context field,
 CTX parameter, 4-9, 5-5
 User control blocks,
 FAB, 4-1
 general, 3-1
 NAM, 7-1
 RAB, 5-1
 XAB, 6-1
 User file directory,
 UFD, B-9
 User file mode, 4-18

User file open, 4-18
 User identification code,
 UIC, B-9
 User record area address field,
 UBF parameter, 5-20
 User record area size field,
 USZ parameter, 5-20
 USZ parameter,
 user record area size field, 5-20

V

V format variable-length records, B-4
 Variable-length records, 4-19
 Variable with fixed-control records, 4-20
 VAX-11 RMS,
 control routines, 15-1
 facilities, 3-1
 functions, 1-1
 routines, 3-2
 VBN,
 virtual block number, B-5
 Virtual block numbers, 5-4, B-5
 VOL parameter,
 relative volume number field, 6-19
 VOL1 label, B-11

W

Wait service,
 \$WAIT macro instruction, 11-30
 WBH record-processing option bit, 5-14
 WCK file-processing option bit, 4-15
 Wild card characters,
 file specifications, 8-5, 8-6
 processing, 13-4, 13-12
 substitution, 7-7
 World user class, 6-8
 Write access rights, 6-8
 Write service,
 \$WRITE macro instruction, 12-7
 Write-behind record-processing option, 5-15
 Write-check file-processing option, 4-16

X

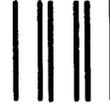
XAB block length field, 6-2
 XAB chain, 6-3

INDEX

XAB parameter,
 extended attribute block
 pointer field, 4-26
XAB type code field, 6-2
XAB types, 6-3
XABALL parameters,
 AID, 6-12
 ALN, 6-13
 ALQ, 6-14
 AOP, 6-14
 BKZ, 6-15
 DEQ, 6-16
 LOC, 6-16
 NXT, 6-3
 VOL, 6-18
XABDAT parameter,
 EDT, 6-5
XABKEY parameters,
 DAN, 6-20
 DFL, 6-21
 DTP, 6-22
 FLG, 6-24
 IAN, 6-26
 IFL, 6-27
 KNM, 6-28
 LAN, 6-28
 NUL, 6-29
 POS, 6-30
 REF, 6-31
 SIZ, 6-32
XABPRO parameters,
 PRO, 6-8
 UIC, 6-10
XAB,
 extended attribute block, 6-1

Do Not Tear - Fold Here and Tape

digital

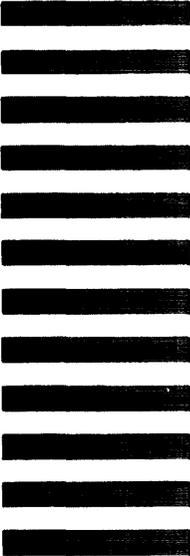


No Postage
Necessary
if Mailed in the
United States

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO.33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

BSSG PUBLICATIONS TW/A14
DIGITAL EQUIPMENT CORPORATION
1925 ANDOVER STREET
TEWKSBURY, MASSACHUSETTS 01876



Do Not Tear - Fold Here