**Systems**

# OS/VS1 Data Management Services Guide

**VS1** Release 6

IBM

# PREFACE

This book describes all IBM data management except for VSAM (virtual storage access method) and specialized applications such as the time sharing option (TSO), graphics, teleprocessing, optical character readers, optical reader-sorters, and magnetic character readers. These specialized applications are described in separate publications that are listed in *IBM System/360 and System/370 Bibliography*, GA22-6822. To learn about VSAM or to write programs that create and process VSAM data sets, refer to:

- *Planning for Enhanced VSAM Under OS/VS*, GC26-3842, which introduces VSAM and describes its concepts and functions.

- *OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide*, GC26-3838, which describes how to create VSAM data sets and code the macro instructions required to process them.

- *OS/VS1 Access Method Services*, GC26-3840, describes the service program commands used to manipulate VSAM data sets.

- *OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications*, GC26-3819, which describes applications not required in the normal use of VSAM.

If you know how to write assembler-language programs and use job control statements, you can use this book and *OS/VS1 Data Management Macro Instructions*, GC26-3872, to write programs that create and process data sets. To use this book you must have basic knowledge of the operating system as contained in *OS/VS1 Planning and Use Guide*, GC24-5090; of assember language as described in *OS/VS-DOS/VS-VM/370 Assembler Language*, GC33-4010; and of job control language (JCL) as explained in *OS/VS1 JCL Services*, GC24-5100, and *OS/VS1 JCL Reference*, GC24-5099.

This book has three parts:

"Part 1: Introduction to Data Management" introduces you to the characteristics of data sets, how you name them, how the system catalogs them, and how you format the records in them. The format of tracks on a direct-access storage device is explained briefly.

Part 1 also describes the data control block (DCB) and the information it supplies to the operating system. Special processing routines that you specify in the DCB macro instruction are also explained in this section.

In "Part 2: Data Management Processing Procedures" there is an explanation of data-processing techniques that includes the macro instructions for the queued access technique and the basic access technique and the macro instructions for analyzing input and output errors. The section on data-processing techniques also tells how to select an access method and how to begin and end processing of a data set.

The section "Buffer Acquisition and Control" in Part 2 explains three different methods you can use to obtain buffers and the macro instructions you use with each method. This section also describes ways to control buffers: simple buffering and exchange buffering for the queued access technique, direct buffering and dynamic buffering for the basic access technique. In addition, for the queued access technique, there is an explanation of the four modes of moving the records in virtual storage: move mode, data mode, locate mode, and substitute mode. Macro instructions for controlling buffers are described here, too.

The next four sections of Part 2 concern processing data sets of four different types: a sequential data set, a partitioned data set, an indexed sequential data set, and a direct

data set. They explain the organization of the data sets and the macro instructions used to process them. In the examples the macro instructions are coded in just enough detail to make the examples clear. For a complete description of the operands and options available, see *OS/VS1 Data Management Macro Instructions,* GC26-3793.

"Part 3: Data Set Disposition and Space Allocation" tells you how to figure the amount of space you need for a data set on a direct-access storage device and how to request that space in your JCL DD statement. You are given special directions for allocating space for a partitioned data set and an indexed sequential data set. Part 3 also tells how to indicate in the JCL DD statement the status of the data set at the beginning of and during processing and how to indicate what you want the system to do with the data set when processing has terminated. You also are told how to use the DD statement to route the data set to a system output writer, to concatenate data sets, to catalog data sets, and to protect confidential data sets.

Appendix A describes data set labeling. Appendix B explains control characters you can use to control card punches and printers. A glossary of acronyms and abbreviations used in this book and the index follow Appendix B.

The following manuals are referred to in the text.

- *OS/VS Message Library: VS1 System Codes,* GC38-1003
- *OS/VS Message Library: VS1 System Messages,* GC38-1001
- *OS/VS1 Data Management for System Programmers,* GC26-3837
- *OS/VS1 JCL Reference,* GC24-5099
- *OS/VS1 JCL Services,* GC24-5100
- *OS/VS1 Service Aids,* GC28-0665
- *OS/VS1 Supervisor Services and Macro Instructions,* GC24-5103
- *OS/VS1 System Data Areas,* SY28-0605
- *OS/VS1 System Generation Reference,* GC26-3791
- *IBM 3800 Printing Subsystem Programmer's Guide,* GC26-3846
- *IBM 3890 Document Processor Machine and Programming Description,* GA24-3612
- *OS Data Management Services and Macro Instructions for IBM 1419/1275,* GC21-5006
- *OS and OS/VS Programming Support for the IBM 3505 Card Reader and IBM 3525 Card Punch,* GC21-5097
- *OS/VS IBM 3886 Optical Character Reader Model 1 Reference,* GC24-5101
- *OS/VS Mass Storage System (MSS) Planning Guide,* GC35-0011
- *OS/VS Mass Storage System (MSS) Services for Space Management,* GC35-0012
- *OS/VS Tape Labels,* GC26-3795
- *OS/VS Utilities,* GC35-0005

# CONTENTS

# FIGURES

# OS/VS1 SUMMARY OF AMENDMENTS

## Release 6

### *New Device Support*

The IBM 3203 Printer, Model 4, is supported with this release.

### *Open Extend Support*

The EXTEND and OUTINX options are supported for the OPEN macro. These options allow the user to change the disposition of a data set to MOD. In all other ways EXTEND and OUTINX are equivalent to the OUTPUT and OUTIN options, respectively.

These new options will allow users of SAM and ISAM to add records to the end of an existing data set even though DISP=OLD/NEW/MOD/ SHR was specified. In the past, the only way to add records to the end of the data set was to specify DISP=MOD on the DD statement and OUTPUT on the OPEN macro or to specify INOUT on the OPEN macro and read to end-of-file or use the OPEN TYPE=J macro.

Information on the IBM 3203 Printer, Model 4, is provided for planning purposes only until the product is available. For additional information about the 3203 Printer, see *IBM 3203 Printer Component Description and Operator's Guide*, GA33-1515.

## Release 5

### *New Programming Support*

- The IBM 3800 Printing Subsystem is supported with this release. For additional programming information for the 3800 and the IEBIMAGE utility program, see *IBM 3800 Printing Subsystem Programmer's Guide*, GC26-3846. Information on the 3800 is provided for planning purposes only until the product is available.

- The IBM 3350 Direct Access Storage is supported with this release. For additional information on the 3350, see *Introduction to IBM 3350 Direct Access Storage*, GA26-1638. Information on the 3350 is provided for planning purposes only until the product is available.

- The IBM 3344 Direct Access Storage is supported with this release. For additional information on the 3344, see *Reference Manual for IBM 3340 Disk Storage*, GA26-1619. Information on the 3344 is provided for planning purposes only until the product is available.

- Chained scheduling is now supported in pageable and nonpageable storage.

## Release 4

### *New Programming Support*

The IBM 3850 Mass Storage System (MSS) is supported with this release. The MSS virtual volumes are functionally equivalent to the 3330/3333 Disk Storage, Model 1. For information on MSS, see *OS/VS Mass Storage System (MSS) Planning Guide*, GC35-0011. MSS information is provided for planning purposes only until the system is available.

## *Editorial Changes*

- The explanation of the EODAD routine has been expanded.

- An explanation of how the SYNAD routine functions with QISAM load mode has been added.

- A list of restrictions when sharing a direct data set in multitasking mode has been added.

- The section titled "Updating a Sequential Data Set" has been expanded.

- A section titled "Writing a Short Block When Using the BSAM WRITE Macro" has been added.

- An explanation of the capacity record (R0) has been added to the section titled "Creating a Direct Data Set."

# PART 1: INTRODUCTION TO DATA MANAGEMENT

## Data Set Characteristics

The data management programs of the operating system help you achieve maximum efficiency in managing the mass of data associated with the many programs that are processed at your installation by providing systematic and effective means of organizing, identifying, storing, cataloging, and retrieving all data, including programs, processed by the operating system.

Data set storage control, along with an extensive catalog system, makes it possible for you to retrieve data by symbolic name alone, without specifying device types and volume serial numbers. In freeing computer personnel from maintaining involved volume serial number inventory lists of stored data and programs, the catalog reduces manual intervention and the likelihood of human error.

Data sets stored within the cataloging system can be classified according to installation needs. For example, a sales department could classify the data it uses by geographic area, by individual salesman, or by any other logical plan.

The cataloging system also makes it possible for you to classify successive generations or updates of related data. These generations can be given an identical name and subsequently be referred to relative to the current generation. The system automatically maintains a list of the most recent generations.

You can request data from a direct-access volume, a remote terminal, or a tape volume, and data organized sequentially or directly, in essentially the same way. In addition, data management provides:

- Allocation of space on direct-access volumes. Flexibility and efficiency of direct-access devices are improved through greater use of available space.

- Automatic retrieval of data sets by name alone.

- Freedom to defer specifications such as buffer length, block size, and device type until a job is submitted for processing. This permits the creation of programs that are in many ways independent of their operating environment.

Control of confidential data is provided by the data set security part of the operating system. You can prevent unauthorized access to payroll data, sales forecast data, and all other data sets that require special security attention. An individual can use a security-protected data set only after furnishing a predefined password.

Input/output routines are provided to efficiently schedule and control the transfer of data between storage and input/output devices. Routines are available to:

- Read data

- Write data

- Translate data from ASCII (American National Standard Code for Information Interchange) to EBCDIC (Extended Binary Coded Decimal Interchange Code) and back

- Block and deblock records

- Overlap reading, writing, and processing operations

- Read and verify volume and data set labels

- Write data set labels

- Automatically position and reposition volumes

- Detect error conditions and correct them when possible

- Provide exits to user-written error and label routines

OS/VS data management programs also provide for a variety of methods for gaining access to a data set. The methods are based on data set organization and data access technique.

OS/VS data sets can be organized in four ways:

- *Sequential:* Records are placed in physical rather than logical sequence. Given one record, the location of the next record is determined by its physical position in the data set. Sequential organization is used for all magnetic-tape devices, and may be selected for direct-access devices. Punched tape, punched cards, and printed output are sequentially organized.

- *Indexed Sequential:* Records are arranged in sequence, according to a key that is a part of every record, on the tracks of a direct-access volume. An index or set of indexes maintained by the system gives the location of certain principal records. This permits direct as well as sequential access to any record.

- *Direct:* The records within the data set, which must be on a direct-access volume, may be organized in any manner you choose. All space allocated to the data set is available for data records. No space is required for indexes. You specify addresses by which records are stored and retrieved directly.

- *Partitioned:* Independent groups of sequentially organized records, called members, are in direct-access storage. Each member has a simple name stored in a directory that is part of the data set and contains the location of the member's starting point. Partitioned data sets are generally used to store programs. As a result, they are often referred to as libraries.

Requests for input/output operations on data sets through macro instructions employ two techniques: the technique for *queued access* and the technique for *basic access*. Each technique is identified according to its treatment of buffering and synchronization of input and output with processing. The combination of an acess technique and a given data set organization is called an *access method*. In choosing an access method for a data set, therefore, you must consider not only its organization, but also what you need to specify through macro instructions. Also, you may choose a data organization according to the access techniques and processing capabilities available.

The code generated by the macro instructions for both techniques is optionally reenterable depending on the form in which parameters are expressed.

In addition to the access methods provided by the operating system, an elementary access technique called *execute channel program* (EXCP) is also provided. To use this technique, you must establish your own system for organizing, storing, and retrieving data. Its primary advantage is the complete flexibility it allows you in using the computer directly.

An important feature of data management is that much of the detailed information needed to store and retrieve data, such as device type, buffer processing technique, and format of output records need not be supplied until the job is ready to be executed. This device independence permits changes to those specifications to be made without changes in the program. Therefore, you may design and test a program without knowing the exact input/output devices that will be used when it is executed.

Device independence is a feature of both access techniques for processing a sequential data set. To some extent, you determine the degree of device independence achieved.

Many useful device-dependent features are available as part of certain macro instructions, and achieving device independence requires some selectivity in their use.

## Data Set Identification

Any information that is a named, organized collection of logically related records can be classified as a data set. The information is not restricted to a specific type, purpose, or storage medium. A data set may be, for example, a source program, a library of macro instructions, or a file of data records used by a processing program.

Whenever you indicate that a new data set is to be created and placed on auxiliary storage, you (or the operating system) must give the data set a name. The data set name identifies a group of records as a data set. All data sets recognized by name (referred to without volume identification) and all data sets residing on a given volume must be distinguished from one another by unique names. To assist in this, the system provides a means of qualifying data set names.

A data set name is one simple name or a series of simple names joined together so that each represents a level of qualification. For example, the data set name DEPT58.SMITH.DATA3 is composed of three simple names. Proceeding from the left, each simple name is a category within which the next simple name is a subcategory.

Each simple name consists of from 1 to 8 alphameric characters, the first of which must be alphabetic. The special character period (.) separates simple names from each other. Including all simple names and periods, the length of the data set name must not exceed 44 characters. Thus, a maximum of 22 simple names can make up a data set name.

To permit different executions of a program to process different data sets without program reassembly, the data set is not referred to by name in the processing program. When the program is executed, the data set name and other pertinent information (such as unit type and volume serial number) are specified in a job control statement called the *data definition (DD)* statement. To gain access to the data set during processing, reference is made to a *data control block (DCB)* associated with the name of the DD statement. Space for a data control block, which specifies the particular data set to be used, is reserved by a DCB macro instruction when your program is assembled.

## Data Set Storage

System/370 provides a variety of devices for collecting, storing, and distributing data. Despite the variety, the devices have many common characteristics. The generic term *volume* is used to refer to a standard unit of auxiliary storage. A volume may be a reel of magnetic tape, a disk pack, or a drum.

Each data set stored on a volume has its name, location, organization, and other control information stored in the *data set label* or *volume table of contents* (for direct-access volumes only). Thus, when the name of the data set and the volume on which it is stored are made known to the operating system, a complete description of the data set, including its location on the volume, can be retrieved. Then, the data itself can be retrieved, or new data added to the data set.

Various groups of labels are used to identify magnetic-tape and direct-access volumes, as well as the data sets they contain. Magnetic-tape volumes can have standard or nonstandard labels, or they can be unlabeled. Direct-access volumes must use standard labels. Standard labels include a volume label, a data set label for each data set, and optional user labels.

Keeping track of the volume on which a particular data set resides can be a burden and a source of error. To alleviate this problem, the system provides for automatic cataloging of data sets. The system can retrieve a cataloged data set if given only the name of the

data set. If the name is qualified, each qualifier corresponds to one of the indexes in the catalog. For example, the system finds the data set DEPT58.SMITH.DATA3 by searching a master index to determine the location of the index name DEPT58, by searching that index to find the location of the index SMITH, and by searching that index for DATA3 to find the identification of the volume containing the data set.

By use of the catalog, collections of data sets related by a common external name and the time sequence in which they were cataloged (their generation) can be identified; they are called *generation data groups*. For example, a data set name LAB.PAYROLL(0) refers to the most recent data set of the group; LAB.PAYROLL(-1) refers to the second most recent data set, etc. The same data set names can be used repeatedly with no requirement to keep track of the volume serial numbers used.

## Direct-Access Volumes

Direct-access volumes are used to store executable programs, including the operating system itself. Direct-access storage is also used for data and for temporary working storage. One direct-access storage volume may be used for many different data sets, and space on it may be reallocated and reused. A volume table of contents (VTOC) is used to account for each data set and available space on the volume.

Each direct-access volume is identified by a volume label, which is stored in track 0 of cylinder 0. You may specify up to seven additional labels, located after the standard volume label, for further identification.

The VTOC is a data set consisting of *data set control blocks (DSCBs)* that describe the contents of the direct-access volume. The VTOC can contain seven kinds of DSCBs, each with a different purpose and a different format number. *OS/VS1 System Data Areas* describes the seven kinds of DSCBs, their purposes, and their formats.

Each direct-access volume is initialized by a utility program before being used on the system. The initialization program generates the volume label and constructs the table of contents. For additional information on direct-access labels, see "Appendix A: Direct-Access Labels."

When a data set is to be stored on a direct-access volume, you must supply the operating system with the amount of space to be allocated to the data set, expressed in blocks, tracks, or cylinders. Space allocation can be independent of device type if the request is expressed in blocks. If the request is made in tracks or cylinders, you must be aware of such device considerations as cylinder capacity and track size.

## Magnetic-Tape Volumes

Because data sets on magnetic-tape devices must be organized sequentially, the operating system does not require space allocation procedures comparable to those for direct-access devices. When a new data set is to be placed on a magnetic-tape volume, you must specify the data set sequence number if it is not the first data set on the reel. The operating system positions a volume with IBM standard labels, American National Standard labels, or no labels so that the data set can be read or written. If the data set has nonstandard labels, you must provide for volume positioning in your nonstandard-label-processing routines. All data sets stored on a given magnetic-tape volume must be recorded in the same density.

When a data set is to be stored on an unlabeled tape volume and you have not specified a volume serial number, the system assigns a serial number to that volume and to any additional volumes required for the data set. Each such volume is assigned a serial number of the form Lxxxyy where xxx indicates the data set sequence number from IPL to IPL and yy indicates the volume sequence number for the data set. If you specify volume serial numbers for unlabeled volumes on which a data set is to be stored, the system assigns volume serial numbers to any additional volumes required. If data sets

residing on unlabeled volumes are to be cataloged or passed, you should specify the volume serial numbers for the volumes required. This will prevent data sets residing on different volumes from being cataloged or passed under identical volume serial numbers. Retrieval of such data sets could result in unpredictable errors.

Each data set and each data set label group on magnetic tape that is to be processed by the operating system must be followed by a tapemark. Tapemarks cannot exist within a data set. When the operating system is used to create a tape with standard labels or no labels, all tapemarks are automatically written. Two tapemarks are written after the last trailer label group on a volume to indicate the last data set on the volume. On an unlabeled volume, the two tapemarks are written after the last data set.

When the operating system is used to create a tape data set with nonstandard labels, the delimiting tapemarks are not written. If the data set is to be retrieved by the operating system, those tapemarks must be written by your nonstandard-label-processing routine. Otherwise, tapemarks are not required after nonstandard labels since positioning of the tape volumes must be handled by installation routines.

For more information on labels for magnetic-tape volumes, refer to *OS/VS Tape Labels*.

The data on magnetic-tape volumes can be in either EBCDIC or ASCII. ASCII is a 7-bit code consisting of 128 characters. It permits data on magnetic tape to be transferred from one computer to another even though the two computers may be products of different manufacturers.

Data management support of ASCII and of American National Standard tape labels is such that data management can translate records on input tapes in ASCII into EBCDIC for internal processing and translate the EBCDIC back into ASCII for output. Records on such input tapes may be sorted into ASCII collating sequence.

## Data Set Record Formats

A data set is composed of a collection of records that normally have some logical relation to one another. The record is the basic unit of information used by a processing program. It might be a single character, all information resulting from a given business transaction, or measurements recorded at a given point in an experiment. Much data processing consists of reading, processing, and writing individual records.

The process of grouping a number of records before writing them on a volume is referred to as *blocking*. A *block* is made up of the data between interrecord gaps (IRGs). Each block can consist of one or more records. Blocking conserves storage space on the volume because it reduces the number of IRGs in the data set. In many cases, blocking also increases processing efficiency by reducing the number of input/output operations required to process a data set.

Records may be in one of four formats: fixed-length (format-F), variable-length for data in EBCDIC (format-V), variable-length for data to be translated to or from ASCII (format-D), or undefined-length (format-U). The main consideration in the selection of a record format is the nature of the data set itself. You must know the type of input your program will receive and the type of output it will produce. Selection of a record format is based on this knowledge, as well as on an understanding of the input/output devices that are used to contain the data set and the access method used to read and write the data records. The record format of a data set is indicated in the data control block according to specifications in the DCB macro instruction, the DD statement, or the data set label.

For ASCII tapes, data can be in format-F, format-D, and format-U with the restrictions noted under "Fixed-Length Records, ASCII tapes," "Variable-Length Records—Format D," and "Undefined-Length Records." When data management reads

records from ASCII tapes, it translates the records into EBCDIC. When data
management writes records onto ASCII tapes, it translates the records into ASCII.
Because you use input records after they are translated and because output records are
translated when you ask data management to write them, you work only with EBCDIC.

**Note:** There is no minimum requirement for block size; however, if a data check occurs
on a magnetic-tape device, any block shorter than 12 bytes in a Read operation or 18
bytes in a Write operation is treated as a noise record and lost. No check for noise is
made unless a data check occurs. The sort/merge program does not accept physical
blocks or logical records shorter than 18 bytes from any device.

## Fixed-Length Records

The size of fixed-length (format-F) records, shown in Figure 1, is constant for all records
in the data set. The number of records within a block is constant for every block in the
data set, unless the data set contains truncated (short) blocks. If the data set contains
unblocked format-F records, one record constitutes one block.

The system automatically performs physical length checking (except for card readers) on
blocked or unblocked format-F records. Allowances are made for truncated blocks.

Format-F records are shown in Figure 1. The optional control character (c), used for
stacker selection or carriage control, may be included in each record to be printed or
punched.
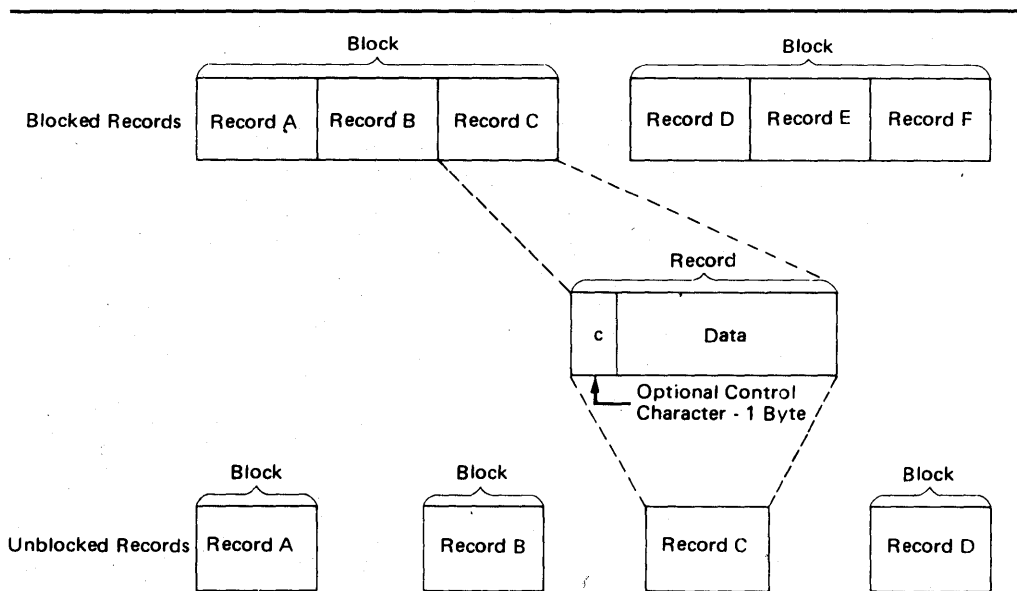


Figure 1.   Fixed-Length Records

**Fixed-Length Records, Standard Format:** During creation of a sequential data set (to be
processed by BSAM or QSAM) with fixed-length records, the RECFM subparameter of
the DCB macro instruction may specify a standard format (RECFM=FS or FBS). A
standard-format data set must conform to the following specifications:

* All records in the data set are format-F records.

* No block except the last block is truncated. (With BSAM you must ensure that this
  specification is met.)

* Every track except the last one contains the same number of blocks.

* Every track except the last one is filled to capacity as determined by the track capacity
  formula established for the device. (These formulas are presented in Part 3 of this
  book under "Allocating Space on Direct-Access Volumes.")

- The data set organization is physical-sequential. A member of a partitioned data set cannot be specified.

A sequential data set with fixed-length records having a standard format can be read more efficiently than a data set that doesn't specify a standard format. This efficiency is possible because the system is able to determine the address of each record to be read because each track contains the same number of blocks.

You should never extend a data set of this type (by coding DISP=MOD) if the last block is truncated, because the extension will cause the data set to contain a truncated block that isn't the last block. This type of data set on magnetic tape should not be read backward, because then the data set would begin with a truncated block. Consequently, you probably won't want to use this type of data set with magnetic tape. If you use one of the basic access techniques with this type of data set, you should not specify that the track overflow feature is to be used with the data set.

Standard format should not be used to read records from a data set that was created using a RECFM other than standard since other record formats may not create the precise format required by standard.

If at any time the characteristics of your data set are altered from the specifications described above, then the data set should no longer be processed with the standard format specification.

**Fixed-Length Records, ASCII Tapes:** For ASCII tapes, format-F records are the same as described above, with two exceptions:

- Control characters, if present, must be American National Standards Institute (ANSI) control characters.

- Records or blocks of records can contain block prefixes.

Figure 2 shows the format of fixed-length records for ASCII tapes and where control characters and block prefixes go if they exist.

The block prefix can vary in length from 0 to 99 bytes but the length must be constant for the data set being processed. For blocked records, the block prefix precedes the first logical record. For unblocked records, the block prefix precedes each logical record.
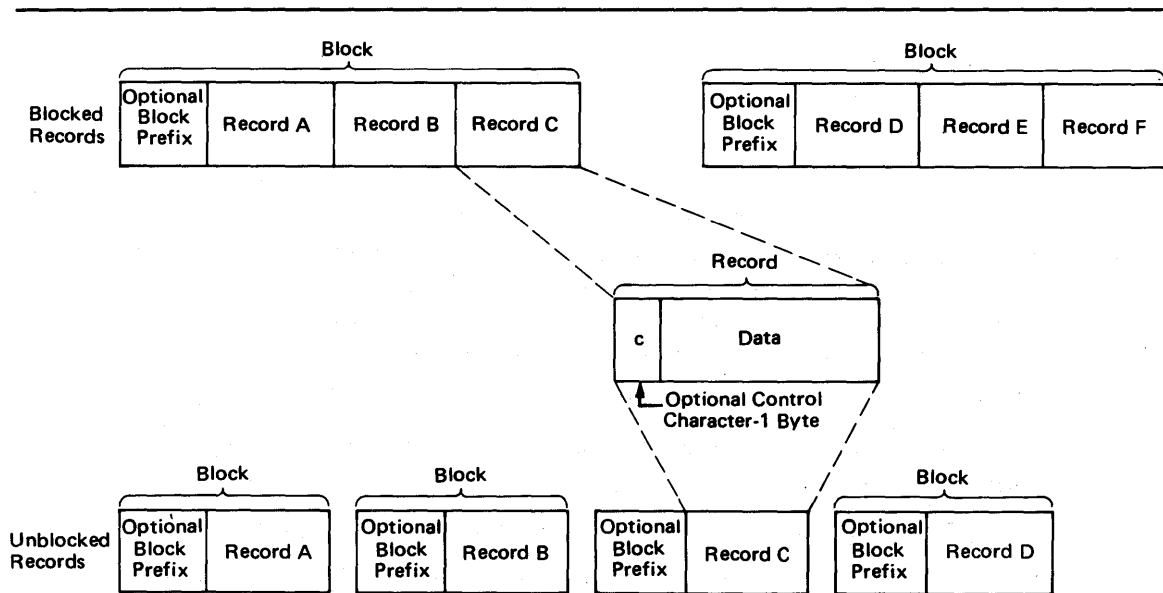


Figure 2. Fixed-Length Records for ASCII Tapes

Using QSAM and BSAM to read records with block prefixes requires that you specify the BUFOFF operand in the DCB. When using QSAM, you cannot read the block prefix on input. When using BSAM, you must account for the block prefix on both input and output. When using either QSAM or BSAM, you must account for the length of the block prefix in the BLKSIZE and BUFL operands of the DCB.

When you use BSAM on output records, the operating system does not recognize a block prefix. Therefore, if you want a block prefix, it must be part of your record. Note that you cannot include block prefixes in QSAM output records.

The block prefix can contain any data you want, but you must avoid using data types such as binary, packed decimal, and floating-point that cannot be translated into ASCII.

For more information about control characters, refer to "Control Character" and to "Appendix B: Control Characters."

## Variable-Length Records

The variable-length record formats are format-V and format-D. Format-V records can be spanned; that is, records can be larger than the blocksize, as described below. Format-D records are used for ASCII tape data sets and cannot be spanned. Figure 3 shows blocked and unblocked variable-length records without spanning.

**Variable-Length Records—Format V:** Format V provides for variable-length records, variable-length record segments, each of which describes its own characteristics, and variable-length blocks of such records or record segments. Except when variable-length track overflow records are specified for volumes on devices with the rotational position sensing feature, the control program performs length checking of the block and uses the record or segment length information in blocking and deblocking. The first 4 bytes of each record, record segment, or block make up a descriptor word containing control information. You must allow for these additional 4 bytes in both your input and output buffers.
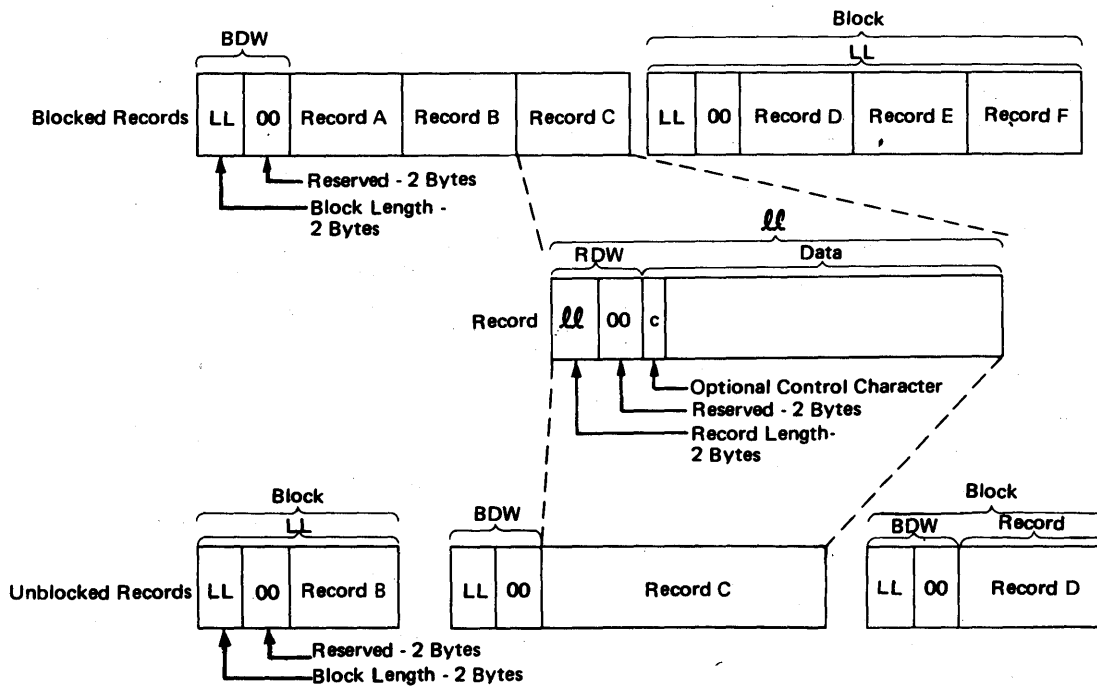


Figure 3. Nonspanned, Variable-Length Records

Block Descriptor Word: A variable-length block consists of a block descriptor word (BDW) followed by one or more logical records or record segments. The block descriptor word is a 4-byte field that describes the block. The first 2 bytes specify the block length ('ll')—4 bytes for the BDW plus the total length of all records or segments within the block. This length can be from 8 to 32,760 bytes or, when you are using WRITE with tape, from 18 to 32,760. The third and fourth bytes are reserved for future system use and must be 0. If the system does your blocking—that is, when you use the queued access technique—the operating system automatically provides the BDW when it writes the data set. If you do your own blocking—that is, when you use the basic access technique—you must supply the BDW.

Record Descriptor Word: A variable-length logical record consists of a record descriptor word (RDW) followed by the data. The record descriptor word is a 4-byte field describing the record. The first 2 bytes contain the length ('ll') of the logical record (including the 4-byte RDW). The length can be from 4 to 32,756. For information about processing a sequential data set, see "Data Format—Device Type Considerations." All bits of the third and fourth bytes must be 0, as other values are used for spanned records. For output, you must provide the RDW except in data mode for spanned records (described under "Buffer Control"). For output in data mode, you must provide the total data length in the physical record length field (DCBPRECL) of the DCB. For input, the operating system provides the RDW except in data mode. In data mode, the system passes the record length to your program in the logical record length field (DCBLRECL) of the DCB. The optional control character (c) may be specified as the fifth byte of each record and must be followed by at least one byte of data (the length in the RDW, in this case, would be six). The RDW and the control character, if specified, are not punched or printed.

*Spanned Variable-Length Records (Sequential Access Method):* The spanning feature of the queued and basic sequential access methods enables you to create and process variable-length logical records that are larger than one physical block and/or to pack blocks with variable-length records by splitting the records into segments so that they can be written into more than one block, as shown in Figure 4.

When spanning is specified for blocked records, the system tries to fill all blocks. For unblocked records, a record larger than blocksize is split and written in two or more blocks, each block containing only one record or record segment. Thus the blocksize may be set to the one that is best for a given device or processing situation. It is not restricted by the maximum record length of a data set. A record may, therefore, span several blocks, and may even span volumes. Note that a logical record spanning three or more volumes cannot be processed in update mode (described under "Buffer Control") by QSAM. A block can contain a combination of records and record segments, but not multiple segments of the same record. When records are added to or deleted from a data set, or when the data set is processed again with different blocksize or record-size parameters, the record segmenting will change.

*Considerations for Processing Spanned Record Data Sets:* When spanned records span volumes, reading errors may occur when using QSAM if a volume which begins with a middle or last segment is mounted first or if an FEOV macro instruction is issued followed by another GET. QSAM cannot begin reading from the middle of the record. The errors include duplicate records, program checks in the user's program, and invalid input from the spanned record data set.

When a spanned record data set is to be opened in UPDAT mode and QSAM is used, a record area must be provided by using the BUILDRCD macro instruction or by specifying BFTEK=A in the DCB.

If you issue the FEOV macro instruction when reading a data set that spans volumes, or if a spanned multivolume data set is opened to other than the first volume, make sure

Figure 4 diagram (Spanned Variable-Length Records):

**BDW**

Block
LL

| LL | | Last Segment of Logical Record A | First Segment of Logical Record B |

| LL | | Intermediate Segment of Logical Record B |

| LL | | Last of Logical Record B | First Segment of Logical Record C |

Reserved - 2 Bytes
Block Length - 2 Bytes

*ll*

SDW   Data

First Segment of Logical Record

| *ll* | | c | |

Optional Control Character
Reserved - 1 Byte
Segment Control Code - 1 Byte (See Figure 5)
Segment Length - 2 bytes

Inter-mediate Segment of Logical Record

*ll*

SDW   Data

| *ll* | | |

Segment Control Code

Last Segment of Logical Record

*ll*

SDW   Data

| *ll* | | |

Segment Control Code

*ll*

RDW        Data Portion of Logical Record B

Logical Record (In User's Work Area)

| *ll* | | c | Data Portion of First Segment | Data Portion of Intermediate Segment | Data Portion of Last Segment |

Optional Control Character
Reserved - 2 Bytes
Record Length - 2 Bytes

Figure 4.   Spanned Variable-Length Records

that each volume begins with the first (or only) segment of a logical record. Input routines cannot begin reading in the middle of a logical record.

Segment Descriptor Word: Each record segment consists of a segment descriptor word (SDW) followed by the data. The segment descriptor word, similar to the record descriptor word, is a 4-byte field that describes the segment. The first 2 bytes contain the length (' $ll$ ') of the segment, including the 4-byte SDW. The length can be from 5 to 32,756 bytes or, when you are using WRITE with tape, from 18 to 32,756 bytes. The third byte of the SDW contains the segment control code, which specifies the relative position of the segment in the logical record. The segment control code is in the rightmost 2 bits of the byte. The segment control codes are shown in Figure 5. The remaining bits of the third byte and all of the fourth byte are reserved for future system use and must be 0.

| Binary Code | Relative Position of Segment |
|---|---|
| 00 | Complete logical record |
| 01 | First segment of a multisegment record |
| 10 | Last segment of a multisegment record |
| 11 | Segment of a multisegment record other than the first or last segment |

Figure 5.   Segment Control Codes

The SDW for the first segment replaces the RDW for the record after the record has been segmented. You or the operating system can build the SDW, depending on which mode of processing is used. In the basic sequential access method, you must create and interpret the spanned records yourself. In the queued sequential access method move mode, complete logical records, including the RDW, are processed in your work area.

GET consolidates segments into logical records and creates the RDW. PUT forms segments as required and creates the SDW for each segment. Data mode is similar to move mode, but allows reference only to the data portion of the logical record in your work area. The logical record length is passed to you through the DCBLRECL field of the data control block. In locate mode, both GET and PUT process one segment at a time. However, in locate mode, if you provide your own record area using the BUILDRCD macro instruction or if you ask the system to provide a record area by specifying BFTEK=A, then GET, PUT, and PUTX process one logical record at a time. (BFTEK=A or the BUILDRCD macro cannot be specified when logical records exceed 32,760 bytes. To process logical records that exceed 32,760 bytes, you must use locate mode and specify LRECL=X in your DCB macro.)

A logical record spanning three or more volumes cannot be processed when the data set is opened for update.

When unit-record devices are used with spanned records, the system assumes that unblocked records are being processed and the block size must be equivalent to the length of one print line or one card. Records that span blocks are written one segment at a time.

SYSIN and SYSOUT Restrictions: Spanned variable-length records cannot be specified for a SYSIN data set. If you're using QSAM to process a SYSOUT data set, move mode (see "Buffer Control") is more efficient than locate mode.

*Null Segments:* A 1 in bit position 0 of the SDW indicates a null segment. A null segment means that there are no more segments in the block. Bits 1-7 of the SDW and the remainder of the block must be binary zeros. A null segment is not an end-of-logical-record delimiter. (You do not have to be concerned about null segments unless you have created a data set using null segments.)

*Spanned Variable-Length Records (Basic Direct Access Method):* The spanning feature of the basic direct access method (BDAM) enables you to create and process variable-length unblocked logical records that are longer than one track. The feature also enables you to pack tracks with variable-length records by splitting the records into segments. These segments can then be written onto more than one track, as shown in Figure 6.

When you specify spanned, unblocked record format for the basic direct access method and when a complete logical record cannot fit on the track, the system tries to fill the track with a record segment. Thus the maximum record length of a data set is not restricted by block size. Furthermore, segmenting records allows a record to span several tracks, with each segment of the record on a different track. However, since the system does not allow a record to span volumes, all segments of a logical record in a direct data set are on the same volume.

**Variable-Length Records—Format D:** For ASCII tapes, variable-length records must be format-D records. Format-D records are the same as format-V records, except:

• Control characters, if present, must be ANSI control characters.

• Records or blocks of records can contain block prefixes.

Figure 7 shows the format of variable-length records for ASCII tapes, where the record descriptor word (RDW) must go, and where block prefixes and control characters can go when they exist.

To specify a block prefix, code the BUFOFF operand in the DCB macro. The block prefix can vary in length from 0 to 99 bytes but its length must remain constant for the data set being processed. For blocked records, the block prefix precedes the first logical record in each block. For unblocked records, the block prefix precedes each logical record. If the block prefix exists, it precedes the RDW.
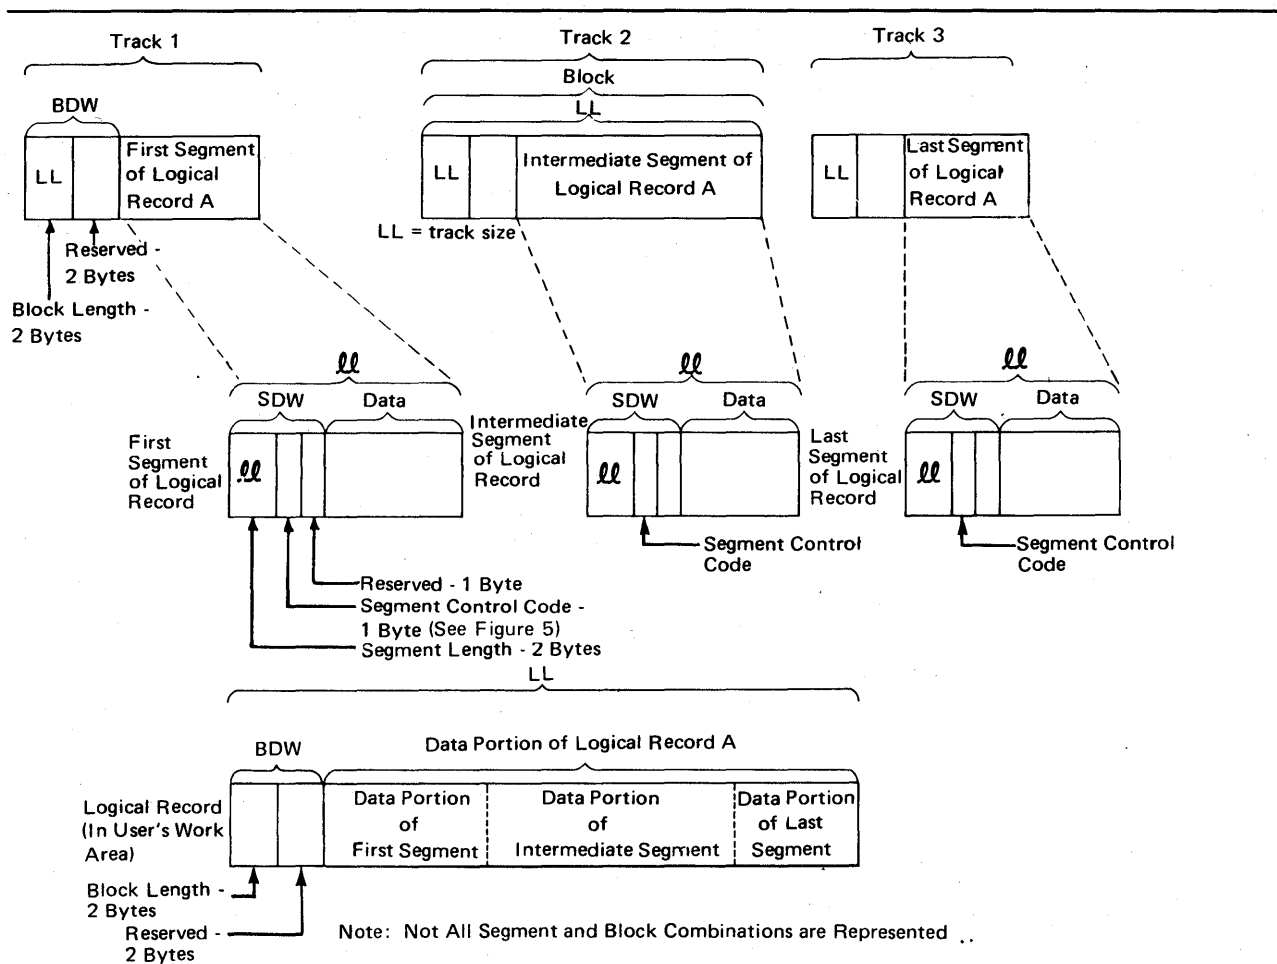
Figure 6. Spanned Variable-Length Records for BDAM Data Sets

To specify that the block prefix is to be treated as a BDW by data management for format-D records on output, code BUFOFF=L as a DCB operand. Your block prefix must be 4 bytes long, and it must contain the length of the block, including the block prefix. The maximum length of a format D, BUFOFF=L block is 9999 because the length (stated in binary by the user) is translated to a four-byte zoned decimal field on the ASCII tape when the tape is written, and is then converted back to a two-byte length field in binary followed by two bytes of zeros when the block is read. If you use QSAM to write records, data management fills in the block prefix for you. If you use BSAM to write records, you must fill in the block prefix yourself. If you are using chained scheduling to read blocked format-D records, coding BUFOFF=*absolute expression* in the DCB is not allowed. Instead, BUFOFF=L is required, because the access method needs binary RDWs and valid end-of-block addresses to unblock the records.

When using QSAM, you cannot read the block prefix on input. When using BSAM, you must account for the block prefix on both input and output. When using either QSAM or BSAM, you must account for the length of the block prefix in the BLKSIZE and BUFL operands.

When you use BSAM on output records, the operating system does not recognize the block prefix. Therefore, if you want a block prefix, it must be part of your record.

The block prefix can contain any data you want, but you must avoid using data types, such as binary, packed decimal, and, floating-point, that cannot be translated into ASCII. For format-D records, the only time the block prefix can contain binary data is when you
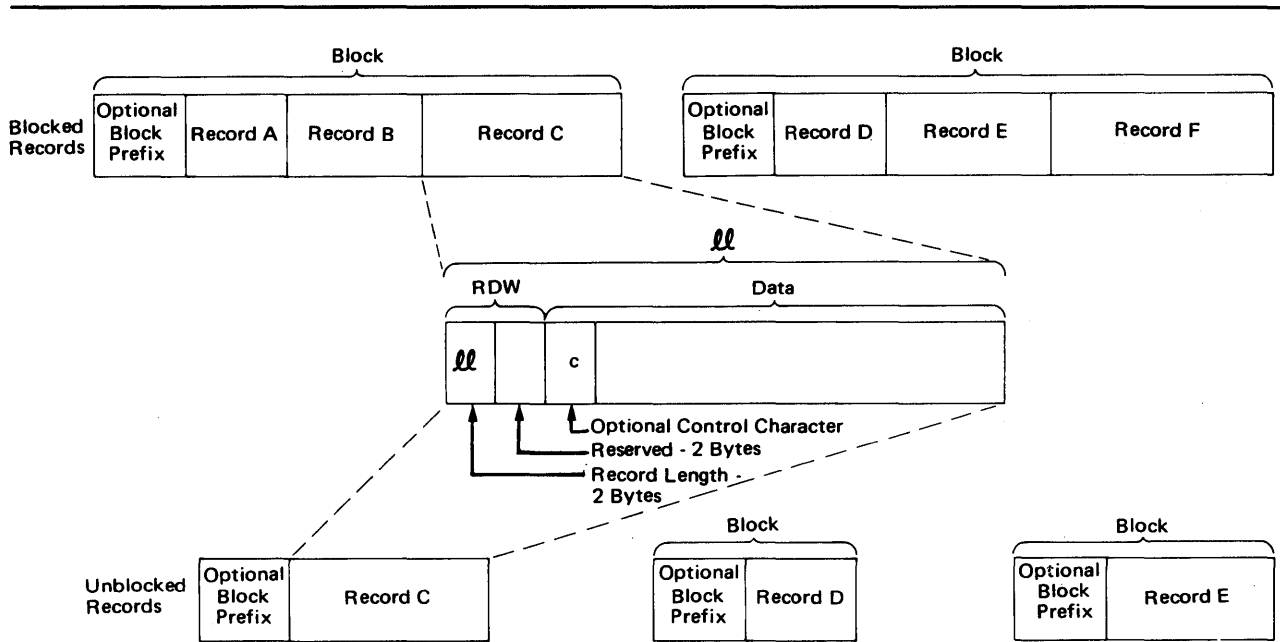
Note: Block prefixes on output records must be 4-bytes long.

Figure 7. Variable-Length Records for ASCII Tapes

have coded BUFOFF=L, which tells data management that the prefix is a BDW. Unlike the block prefix, the RDW must always be in binary.

If you create variable-length records that are shorter than 18 bytes, data management pads each one up to a length of 18 bytes when the records are written onto ASCII tape. The padding character used is the ASCII circumflex.

For more information about control characters, refer to "Control Character" and to "Appendix B: Control Characters."

## Undefined-Length Records

Format U permits processing of records that do not conform to the F or V format. As shown in Figure 8, each block is treated as a record; therefore, deblocking must be performed by your program. The optional control character may be used in the first byte of each record. Because the system does not perform length checking on format-U records, your program may be designed to read less than a complete block into virtual storage.

For ASCII tapes, format-U records are the same as described above, with the two exceptions described for format-F records on ASCII tapes.

Figure 9 shows the format of undefined-length records for ASCII tapes and where a control character and block prefix, if any, go.

For format-U records, the user must specify the record length when issuing the WRITE, PUT, or PUTX macro instruction. No length checking is performed by the system, so no error indication will be given if the specified length does not match the buffer size or physical record size.
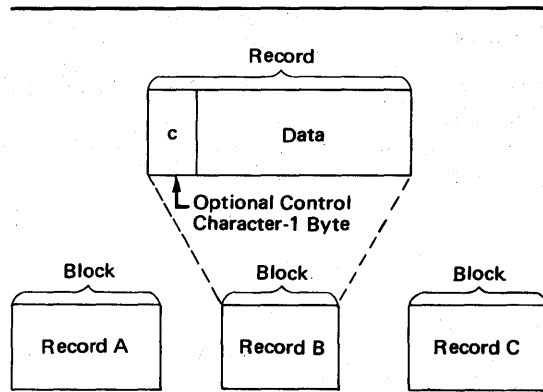
Figure 8. Undefined-Length Records
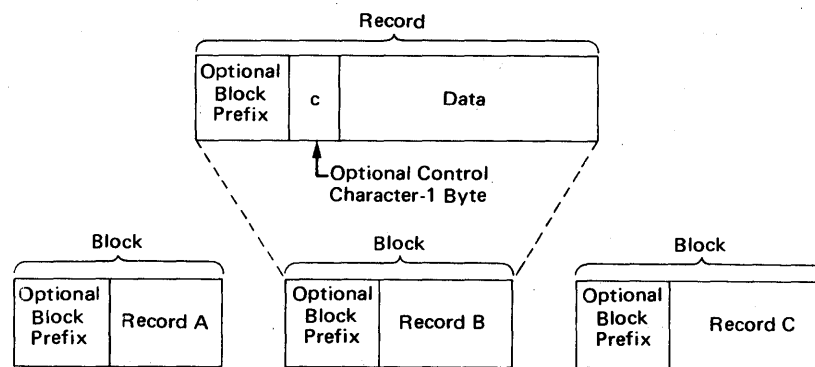


Figure 9. Undefined-Length Records for ASCII Tapes

In update mode, you must issue a GET or READ macro before you issue a PUTX or WRITE macro to a data set on a direct-access device. If you change the record length when you issue the PUTX or WRITE macro, the record will be padded with zeros or truncated to match the length of the record received when the GET or READ macro was issued. No error indication will be given.

**Control Character**

You may specify in the DD statement, the DCB macro instruction, or the data set label that an optional control character is part of each record in the data set. The 1-byte character is used to indicate a carriage control channel when the data set is printed or a stacker bin when the data set is punched. Although the character is a part of the record in storage, it is never printed or punched. For that reason, buffer areas must be large enough to accommodate the character. If the immediate destination of the record is a device, such as disk, that does not recognize the control character, the system assumes that the control character is the first byte of the data portion of the record. If the destination of the record is a printer or punch and you have not indicated the presence of a control character, the system regards the control character as the first byte of data. A list of the control characters is in "Appendix B: Control Characters."

# Direct-Access Device Characteristics

Regardless of organization, data sets created using the operating system can be stored on a direct-access volume. Each block of data has a distinct location and a unique address, making it possible to locate any record without extensive searching. Thus, records can be stored and retrieved either directly or sequentially.

Although direct-access devices differ in physical appearance, capacity, and speed, they are similar in data recording, data checking, data format, and programming. The recording surface of each volume is divided into many concentric *tracks*. The number of tracks and their capacity vary with the device. Each device has some type of *access mechanism*, containing read/write heads that transfer data as the recording surface rotates past them. Only one head at a time can transfer data.

The logical arrangement of related tracks is vertical rather than horizontal. As shown in Figure 10, a cylinder of a 2316 disk pack is composed of 20 tracks, one for each recording surface. Because there are 203 tracks per recording surface, there are 203 vertical cylinders of 20 tracks each. If a data set extends to more than 1 track, it is continued on the next track in the cylinder, not the next track on the same recording surface.



Figure 10. 2316 Disk Pack

## Track Format

Information is recorded on all direct-access volumes in a standard format. In addition to device data, each track contains a track descriptor record (*capacity record* or R0) and data records.

As shown in Figure 11, there are two possible data record formats—count-data and count-key-data—only one of which can be used for a particular data set.

In addition to device data, the count area contains 8 bytes that identify the location of the record by cylinder, head, and record numbers, its key length (0 if no keys are used), and its data length.

If the records are written with keys, the key area (1 to 255 bytes) contains a record key that specifies the data record by part number, account number, sequence number, or some other identifier. In some cases, records are written with keys so that they can be located quickly.

## Track Addressing

Two types of addresses can be used to store and retrieve data on a direct-access volume: actual addresses and relative addresses. The only advantage of using actual addresses is the elimination of time required to convert from relative to actual addresses and vice versa. When sequentially processing a multiple-volume data set, you can refer only to records of the current volume.

**Actual Addresses:** When the system returns the actual address of a record on a direct-access volume to your program, it is in the form MBBCCHHR, where:

M

is a 1-byte binary number specifying the relative location of an entry in a data extent block (DEB). The data extent block is created by the system when the data set is opened. Each extent entry describes a set of consecutive tracks allocated for the data set.

BBCCHH

is three 2-byte binary numbers specifying the cell (bin), cylinder, and head number for the record (its track address). The cylinder and head numbers are recorded in the count area for each record.
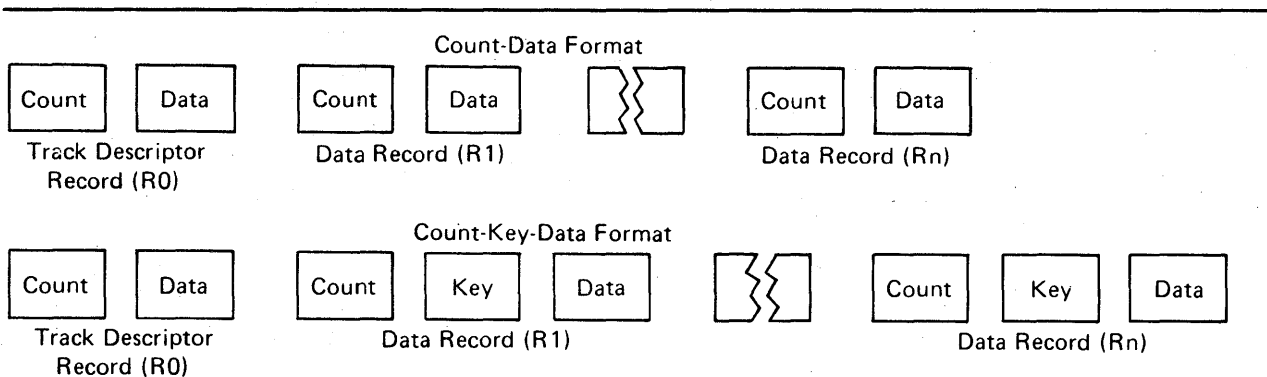


Figure 11. Direct-Access Volume Track Formats

**R**

> is a 1-byte binary number specifying the relative block number on the track. The block number is also recorded in the count area.

If you use actual addresses in your program, the data set must be treated as unmovable.

**Relative Addresses:** Two kinds of relative addresses can be used to refer to records in a direct-access data set: relative block addresses and relative track addresses.

The relative block address is a 3-byte binary number that indicates the position of the block relative to the first block of the data set. Allocation of noncontinuous sets of blocks does not affect the number. The first block of a data set always has a relative block address of 0.

The relative track address has the form TTR, where:

**TT**

> is a 2-byte binary number specifying the position of the track relative to the first track allocated for the data set. The TT for the first track is 0. Allocation of noncontinuous sets of tracks does not affect the number.

**R**

> is a 1-byte binary number specifying the number of the block relative to the first block on the track TT. The R value for the first block of data on a track is 1.

## Track Overflow

If the record overflow feature is available for the direct-access device being used, you can reduce the amount of unused space on the volume by specifying the *track overflow option* in the DD statement or the DCB macro instruction associated with the data set. If the option is used, a block that does not fit on the track is partially written on that track and continued on the next track. (The track onto which the record is continued must be physically next and must be part of the same extent as the track that holds the first part of the record.) Each segment (the portion written on one track) of an overflow block has a count area. The data length field in the count area specifies the length of that segment only. If the block is written with a key, there is only one key area for the block. It is written with the first segment. If the track overflow option is not used, blocks are not split between tracks.

## Write-Validity-Check Option

You can specify the *write-validity-check option* in either the DD statement or the DCB macro instruction. After a record is transferred from main to secondary storage, the system reads the stored record (without data transfer) and, by testing for a data check from the I/O device, verifies that the record was written correctly. This verification requires an additional revolution of the device for each record that was written. Standard error recovery procedures are initiated if an error condition is detected.

# The Data Control Block

You must describe the characteristics of a data set, the volume on which it resides, and its processing requirements before processing can begin. During execution, the descriptive information is made available to the operating system in the *data control block (DCB)*. A DCB is required for each data set and is created in a processing program by a DCB macro instruction.

Primary sources of information to be placed in the data control block are a DCB macro instruction, a data definition (DD) statement, and a data set label. In addition, you can provide or modify some of the information during execution by storing the pertinent data in the appropriate field of the data control block. The specifications needed for input/output operations are supplied during the initialization procedures of the OPEN macro instruction. Therefore, the pertinent data can be provided when your job is to be executed rather than when you write your program (see Figure 12).

When the OPEN macro instruction is executed, the Open routine:

- Completes the data control block

- Loads all necessary access method routines not already in virtual storage

- Initializes data sets by reading or writing labels and control information

- Constructs the necessary system control blocks

Information from a DD statement is stored in the *job file control block (JFCB)* by the operating system. When the job is to be executed, the JFCB is made available to the open routine. The data control block is filled in with information from the DCB macro instruction, the JFCB, or an existing data set label. If more than one source specifies information for a particular field, only one source is used. A DD statement takes precedence over a data set label, and a DCB macro instruction over both. However, you can modify most data control block fields either before the data set is opened or when the operating system returns control to your program (at the data control block open exit). Some fields can be modified during processing.

Figure 13 illustrates the process and the sequence of filling in the data control block from various sources. The primary source is your program, that is, the DCB macro instruction. In general, you should use only those DCB parameters that are needed to ensure correct processing. The other parameters can be filled in when your program is to be executed. When a direct-access data set is opened (or a magnetic tape with standard labels is

| DCB Macro | DD Statement | Data Set Label |
|---|---|---|
| B  F  G  H  J | C    D    I | A        E |

Data Control Block

A  B  C  D  E  F  G  H  I  J
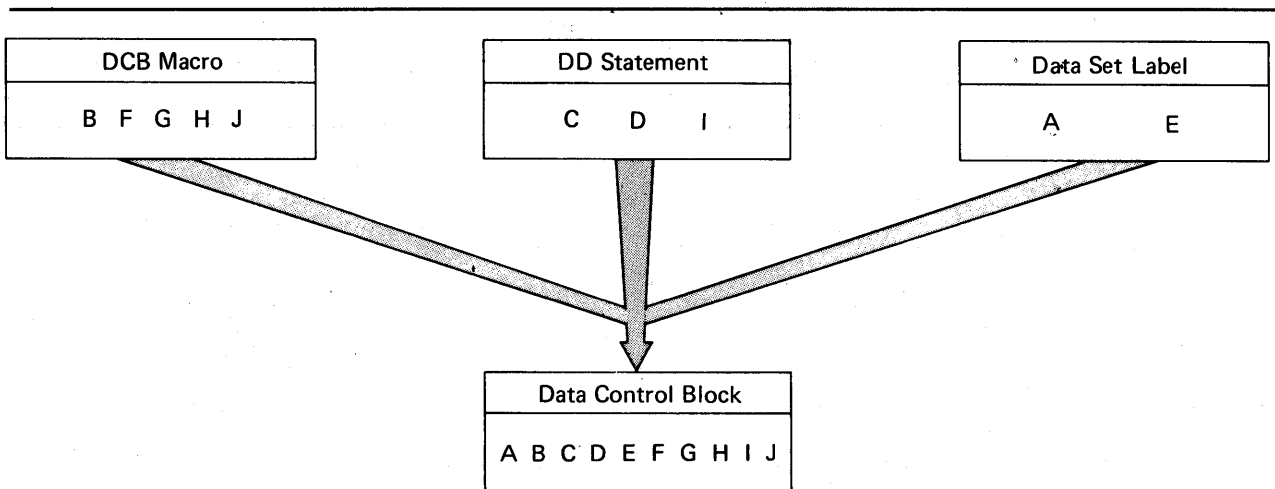
Figure 12. Completing the Data Control Block

opened for INPUT, RDBACK, or INOUT), any field in the JFCB not completed by a
DD statement is filled in from the data set label (if one exists). When opening a magnetic
tape for output, the tape label is assumed not to exist or to apply to the current data set
unless you specify DISP=MOD and a volume serial number in the volume parameter of
the DD statement. Any field not completed in the DCB is filled in from the JFCB. Fields
in the DCB can then be completed or modified by your own DCB exit routine. Then all
DCB fields are unconditionally merged into corresponding JFCB fields if your data set is
| opened for output (OUTPUT, OUTIN, EXTEND, or OUTINX is specified in the OPEN
macro instruction). The DSORG field is not merged unless this field contains zeros in the
JFCB. If your data set is opened for input (INPUT, INOUT, RDBACK, or UPDAT is
specified in the OPEN macro instruction), the DCB fields are not merged unless the
corresponding JFCB fields contain zeros.
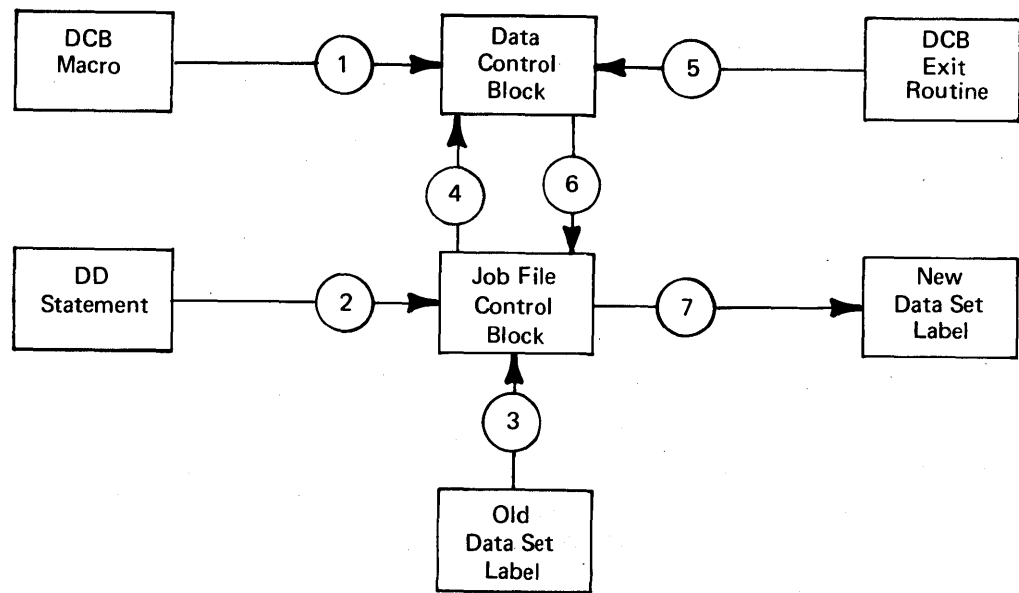


Figure 13. Sources and Sequence of Operations for Completing the Data Control Block

When the data set is closed, the data control block is restored to the condition it had
before the data set was opened (the buffer pool is not freed). The open and close
routines also use the updated JFCB to write the data set labels for output data sets. If the
data set is not closed when your program terminates, the operating system will close it
automatically. Note, however, that the system cannot automatically close any open data
sets after the normal termination of a program that was brought into virtual storage by
the loader. Therefore, loaded programs must include CLOSE macro instructions for all
opened data sets.

## Data Set Description

For each data set you are going to process, there must be a corresponding DCB and DD
statement. The characteristics of the data set and device-dependent information can be
supplied by either source. In addition, the DD statement must supply data set
identification, device characteristics, space allocation requests, and related information as
specified in *OS/VS1 JCL Reference*. You establish the logical connection between a
DCB and a DD statement by specifying the name of the DD statement in the DDNAME
field of the DCB macro instruction, or by completing the field yourself before opening
the data set.

Once the data set characteristics have been specified in the DCB macro instruction, they can be changed only by modification of the DCB during execution. The fields of the DCB discussed below are common to most data organizations and access techniques.

**Data Set Organization (DSORG):** specifies the organization of the data set as physical sequential (PS), indexed sequential (IS), partitioned (PO), or direct (DA). If the data set contains absolute rather than relative addresses, you must mark it as unmovable by adding a U to the DSORG parameter (for example, by coding DSORG=PSU). You must specify the data set organization in the DCB macro instruction. When creating or processing an indexed sequential organization data set or creating a direct data set, you must also specify DSORG in the DD statement.

**Record Format (RECFM):** specifies the characteristics of the records in the data set as fixed-length (F), variable-length (V), or undefined-length (U). Blocked records are specified as FB or VB. You may also specify the records as fixed-length standard by using FS or FBS. You can request track overflow for records other than standard format by adding a T to the RECFM parameter (for example, by coding FBT).

**Record Length (LRECL):** specifies the length, in bytes, of each record in the data set. If the records are of variable length, the maximum record length must be specified. For input, the field should be omitted for format-U records.

**Blocksize (BLKSIZE):** specifies the maximum length, in bytes, of a block. If the records are of format F, the blocksize must be an integral multiple of the record length except for SYSOUT data sets. (See "Routing Data Through the System Input and Output Streams" in Part 3 of this book.) If the records are of format V, the blocksize specified must be the maximum blocksize. If records are unblocked, the blocksize must be 4 bytes greater than the record length (LRECL). When spanned variable-length records are specified, the blocksize is independent of the record length.

**Key Length (KEYLEN):** specifies the length (0-255) in bytes of an optional key which precedes each block on a direct-access device. The value of KEYLEN is not included in BLKSIZE or LRECL but must be included in BUFL if buffer length is specified. Thus, BUFL=KEYLEN+BLKSIZE.

Each of the data set description fields of the data control block, except as noted for data set organization, can be specified when your job is to be executed. In addition, data set identification and disposition, as well as device characteristics, can be specified at that time. The parameters of the DD statement discussed below are common to most data set organizations and devices.

**Data Definition Name (DDNAME):** is the name of the DD statement and connects the DD statement to the data control block that specifies the same DDNAME.

**Data Set Name (DSNAME):** specifies the name of a newly defined data set, or refers to a previously defined data set.

**Data Control Block (DCB):** provides, by means of subparameters, information to be used to complete those fields of the data control block that were not specified in the DCB macro instruction. This parameter cannot be used to modify a data control block.

**Channel Separation and Affinity (SEP/AFF):** requests that specified data sets use different channels during input/output operations.

**Input/Output Device (UNIT):** specifies the number and type of I/O devices to be allocated for use by the data set.

**Space Allocation (SPACE):** designates the amount of space on a direct-access volume that should be allocated for the data set. Unused space can be released when your job is finished.

**Volume Identification (VOLUME):** identifies the particular volume or volumes, or the number of volumes, to be assigned to the data set, or the volumes on which existing data sets reside.

**Data Set Label (LABEL):** indicates the type and contents of the label or labels associated with the data set. The operating system verifies standard labels. Standard labels include those specified in the DD statement as SL (standard labels), SUL (standard user labels), AL (American National Standard labels), and AUL (American National Standard user labels). Nonstandard labels (NSL) can be specified only if your installation has incorporated into the operating system routines to write and process nonstandard labels.

**Data Set Disposition (DISP):** describes the status of a data set and indicates what is to be done with it at the end of the job step.

## Processing Program Description

The operating system requires several types of processing information to ensure proper control of your input/output operations. The forms of macro instructions in the program, buffering requirements, and the addresses of your special processing routines must be specified during either the assembly or the execution of your program. The DCB parameters specifying buffer requirements are discussed in "Buffer Acquisition and Control."

Because macro instructions are expanded during the assembly of your program, you must supply the macro instruction forms that are to be used in processing each data set in the associated DCB macro instruction. You can supply buffering requirements and related information in the DCB macro instruction, the DD statement, or by storing the pertinent data in the appropriate field of the data control block before the end of your DCB exit routine. If the addresses of special processing routines are omitted from the DCB macro instruction, you must complete them in the DCB before opening the data set.

### Macro Instruction Form (MACRF)

The MACRF parameter of the DCB macro instruction specifies not only the macro instructions used in your program, but also the processing mode as discussed in the section "Buffer Control." The organization of your data set, the macro instruction form, and the processing mode determine which of the data access routines will be used during execution.

### Exits to Special Processing Routines

The DCB macro instruction can be used to identify the location of:

- A routine that performs end-of-data procedures

- A routine that supplements the operating system's error recovery routine

- A list that contains addresses of special exit routines

The exit addresses can be specified in the DCB macro instruction or you can complete the DCB fields before opening the data set. Figure 14 summarizes the exits that you can specify either explicitly in the DCB, or implicitly by specifying the address of an exit list in the DCB.

**End-of-Data-Set Exit Routine (EODAD):** The EODAD parameter of the DCB macro instruction specifies the address of your end-of-data routine, which may perform any final processing on an input data set. This routine is entered when an FEOV macro is issued or when a CHECK or GET macro is issued and there are no more records or blocks to be retrieved. (On a READ request, this routine is entered when you issue a CHECK macro instruction to check for completion of the read operation. For a BSAM

| Exit Routine | When Available | Where Specified |
|---|---|---|
| End-of-Data-Set | When no more sequential records or blocks are available | EODAD operand |
| Error Analysis | After an uncorrectable input/output error | SYNAD operand |
| Standard User Label (physical sequential or direct organization) | When opening, closing, or reaching the end of a data set, and when changing volumes | EXLST operand and exit list |
| DCB Open | When opening a data set | EXLST operand and exit list |
| JFCBE | When opening a data set for the 3800 | EXLST operand and exit list |
| End-of-Volume | When changing volumes | EXLST operand and exit list |
| Block Count | After unequal block count comparison by end-of-volume routine | EXLST operand and exit list |
| FCB Image | When opening a data set or issuing a SETPRT macro | EXLST operand and exit list |
| DCB ABEND | When an ABEND condition occurs in Open, Close, or end-of-volume routine. | EXLST operand and exit list |

Figure 14. Data Management Exit Routines

data set that is opened for UPDAT, this routine is entered at the end of each volume. This allows you to issue WRITE macros before an FEOV macro is issued.)

The EODAD routine is not a subroutine, but rather a continuation of the routine which issued the CHECK, GET, or FEOV macro instruction. Once in your EODAD routine, you can continue normal processing, such as reposition and resume processing of the data set, close the data set, or process another data set.

For BSAM, you must first reposition the data set that reached end-of-data if you wish to issue a BSP, READ, or WRITE macro instruction. You can reposition your data set by issuing a CLOSE TYPE=T macro instruction. If a READ macro is issued before the data set is repositioned, unpredictable results will occur.

For BPAM, you may reposition the data set by issuing a FIND or POINT macro instruction. (CLOSE TYPE=T with BPAM results in a no operation performed.)

For QISAM, you can continue processing the input data set that reached end-of-data by first issuing an ESETL macro to end the sequential retrieval, then issuing a SETL macro to set the lower limit of sequential retrieval. You can then issue GET macros to the data set.

Your task will be abnormally terminated under either of the following conditions:

• No exit routine is provided.

• A GET macro instruction is issued in the EODAD routine to the DCB which caused this routine to be entered (unless the access method is QISAM).

When control is passed to the EODAD routine, the registers contain the following information:

| Register | Contents |
|---|---|
| 0-1 | Reserved |
| 2-13 | Contents before execution of CHECK, GET, or FEOV macro instruction |
| 14 | Address of the instruction after the last issued GET, CHECK, or FEOV macro instruction |
| 15 | Reserved |

**Synchronous Error Routine Exit (SYNAD):** The SYNAD parameter of the DCB macro instruction specifies the address of an error routine that is to be given control when an input/output error occurs. This routine can be used to analyze exceptional conditions or uncorrectable errors. The block being read or written can be accepted or skipped, or processing can be terminated.

If an input/output error occurs during data transmission, standard error recovery procedures, provided by the operating system, attempt to correct the error before returning control to your program. An uncorrectable error usually causes an abnormal termination of the task. However, if you specify in the DCB macro instruction the address of an error analysis routine (called a SYNAD routine), the routine is given control in the event of an uncorrectable error.

You can write a SYNAD routine to determine the cause and type of error that occurred by examining:

- The contents of the general registers

- The data event control block (discussed in Part 2 under "Basic Access Technique")

- The exceptional condition code

- The standard status and sense indicators

You can use the SYNADAF macro instruction to perform this analysis automatically. This macro instruction produces an error message that can be printed by a subsequent PUT or WRITE macro instruction.

After completing the analysis, you can return control to the operating system or close the data set. If you close the data set, note that you may not use the temporary close (CLOSE TYPE=T) option in the SYNAD routine. To continue processing the same data set, you must first return control to the control program by a RETURN macro instruction. The control program then transfers control to your processing program, subject to the conditions described below. In no case should you attempt to reread or rewrite the record, because the system has already attempted to recover from the error.

When you are using GET and PUT to process a sequential data set, the operating system provides three automatic error options (EROPT) to be used if there is no SYNAD routine or if you want to return control to your program from the SYNAD routine:

- ACC    accept the erroneous block

- SKP    skip the erroneous block

- ABE    abnormally terminate the task

These options are applicable only to data errors, as control errors result in abnormal termination of the task. Data errors affect only the validity of a block of data. Control errors affect information or operations necessary for continued processing of the data set. These options are not applicable to output errors, except output errors on the printer. When chained scheduling is used, the SKP option is not available, and ACC is assumed if SKP is coded. If the EROPT and SYNAD fields are not completed, ABE is assumed.

When you use READ and WRITE macro instructions, errors are detected when you issue a CHECK macro instruction. If you are processing a direct or sequential data set and you return to the control program from your SYNAD routine, the operating system assumes that you have accepted the bad record. If you are creating a direct data set and you return to the control program from your SYNAD routine, your task is abnormally terminated. In the case of processing a direct data set, the return should be made to the control program via register 14 in order to make a control block (the IOB) available for reuse in a subsequent READ or WRITE macro instruction.

For a detailed description of the register contents upon entry to your SYNAD routine, refer to the tables in *OS/VS1 Data Management Macro Instructions*. The tables there describe register contents for programs using QISAM, BISAM, BDAM, BPAM, BSAM, and QSAM.

Your SYNAD routine can end by branching to another routine in your program, such as a routine that closes the data set. It can also end by returning control to the control program, which then returns control to the next sequential instruction (after the macro) in your program. If your routine returns control, the conventions for saving and restoring register contents are as follows:

- The SYNAD routine must preserve the contents of registers 13 and 14. If required by the logic of your program, the routine must also preserve the contents of registers 2 through 12. Upon return to your program, the contents of registers 2 through 12 will be the same as upon return to the control program from the SYNAD routine.

- The SYNAD routine must not use the save area whose address is in register 13, because this area is used by the control program. If the routine saves and restores register contents, it must provide its own save area.

- If the SYNAD routine calls another routine or issues supervisor or data management macro instructions, it must provide its own save area or issue a SYNADAF macro instruction. The SYNADAF macro instruction provides a save area for its own use, and then makes this area available to the SYNAD routine. Such a save area must be removed from the save area chain by a SYNADRLS macro instruction before control is returned to the control program.

When you use QSAM to read and translate paper-tape characters, your SYNAD routine receives control when you request the record preceding the record in error. Before giving control to your SYNAD routine, the system translates the requested record into your buffer.

For example, suppose that you are using QSAM to read and translate a paper-tape data set and that you have specified, in your DCB, SYNAD=(address) and EROPT=ACC. Suppose also that the third record of the data set has a parity error. When you issue a GET request for the second record, the system translates that record into your buffer and, as a result of the error in the third record, passes control to your SYNAD routine. Because you specified the accept option, the system returns control to your program after your SYNAD error analysis routine completes its processing. When you issue a GET request for the third record, all characters other than the erroneous one are translated into your buffer; the erroneous character is moved, in normal sequence, into your buffer without translation.

If the error analysis routine receives control from the Close routine when indexed sequential data sets are being created (the DCB is opened for QISAM load mode), bit 3 of the IOBFLAGS field in the load mode buffer control table (IOBBCT) is set to one. The DCBWKPT6 field in the DCB contains an address of a list of work area pointers (ISLVPTRS). The pointer to the IOBBCT is at offset 8 in this list as shown in the following diagram:



If the error analysis routine receives control from the Close routine when indexed sequential data sets are being processed using QISAM scan mode, bit 2 of the DCB field DCBEXCD2 is set to one.

**Exit List (EXLST):** The EXLST parameter of the DCB macro instruction specifies the address of a list that contains the addresses of special processing routines, a forms control buffer (FCB) image, or a user totaling area. An exit list must be created if user label, data control block, end-of-volume, block count, JFCBE, or DCB ABEND exits are used, or if a PDAB macro or FCB image is defined in the processing program.

The exit list is constructed of 4-byte entries that must be aligned on fullword boundaries. Each exit list entry is identified by a code in the high-order byte, and the address of the routine, image, or area is specified in the 3 low-order bytes. Codes and addresses for the exit list entries are shown in Figure 15.

You can activate or deactivate any entry in the list by placing the required code in the high-order byte. Care must be taken, however, not to destroy the last entry indication. The operating system routines scan the list from top to bottom, and the first active entry found with the proper code is selected.

You can shorten the list during execution by setting the high-order bit to 1, and extend it by setting the high-order bit to 0.

When control is passed to an exit routine, the registers contain the following information:

| Register | Contents |
|---|---|
| 0 | Variable; see exit routine description. |
| 1 | The three, low-order bytes contain the address of DCB currently being processed, except when user-label exits (X'01'-'04'), user totaling exit (X'0A'), or DCB ABEND exit (X'11') is taken, when register 1 contains the address of a parameter list. The contents of the parameter list are described in each exit routine description. |
| 2-13 | Contents before execution of the macro instruction. |
| 14 | Return address (must not be altered by the exit routine). |
| 15 | Address of exit routine entry point. |

The conventions for saving and restoring register contents are as follows:

• The exit routine must preserve the contents of register 14. It need not preserve the contents of other registers. The control program restores the contents of registers 2-13 before returning control to your program.

• The exit routine must not use the save area whose address is in register 13, because

| Entry Type | Hexadecimal Code | 3-Byte Address—Purpose |
|---|---|---|
| Inactive entry | 00 | Ignore the entry; it is not active. |
| Input header label | 01 | Process a user input header label. |
| Output header label | 02 | Create a user output header label. |
| Input trailer label | 03 | Process a user input trailer label. |
| Output trailer label | 04 | Create a user output trailer label. |
| Data control block exit | 05 | Take a data control block exit. |
| End-of-volume | 06 | Take an end-of-volume exit. |
| User totaling | 0A | Address of beginning of user's totaling area. |
| Block count exit | 0B | Take a block-count-unequal exit. |
| Defer input trailer label | 0C | Defer processing of a user input trailer label from end-of-data until closing. |
| Defer nonstandard input trailer label | 0D | Defer processing a nonstandard input trailer label on magnetic tape unit from end-of-data until closing (no exit routine address). |
| FCB image | 10 | Define an FCB image. |
| DCB ABEND exit | 11 | Examine the ABEND condition and select one of several options. |
| QSAM parallel input | 12 | Address of the PDAB for which this DCB is a member. |
| JFCBE | 15 | Take an exit during open to allow user to examine JCL-specified setup requirements for a 3800 printer. |
| Last entry | 80 | Treat this entry as last entry in list. This code can be specified with any of the above but must always be specified with the last entry. |

Figure 15. Format and Contents of an Exit List

this area is used by the control program. If the exit routine calls another routine or issues supervisor or data management macro instructions, it must provide the address of a new save area in register 13.

*Standard User Label Exit:* When you create a data set with physical sequential or direct organization, you can provide routines to create your own data set labels. You can also provide routines to verify these labels when you use the data set as input. Each label is 80 characters long with the first 4 characters UHL1,UHL2,...,UHL8 for a header label or UTL1,UTL2,...,UTL8 for a trailer label. User labels are not allowed on indexed sequential data sets.

The physical location of the labels on the data set depends on the data set organization. For direct (BDAM) data sets, user labels are placed on a separate user label track in the first volume. User label exits are taken only during execution of the open and close routines. Thus you may create or examine up to eight user header labels only during execution of open and up to eight trailer labels only during execution of close. Since the trailer labels are on the same track as the header labels, the first volume of the data set must be mounted when the data set is closed.

For physical sequential (BSAM or QSAM) data sets, you may create or examine up to eight header labels and eight trailer labels on each volume of the data set. For ASCII tape data sets, you may create an unlimited number of user header and trailer labels. The user label exits are taken during open, close, and end-of-volume processing.

To create or verify labels, you must specify the addresses of your label exit routines in an exit list as shown in Figure 15. Thus you may have separate routines for creating or verifying header and trailer label groups. Care must be taken if a magnetic tape is read

backward, since the trailer label group is processed as header labels and the header label group is processed as trailer labels.

When your routine receives control, the contents of register 0 are unpredictable. Register 1 contains the address of a parameter list. The contents of registers 2-13 are the same as when the macro instruction was issued. However, if your program does not issue the CLOSE macro instruction, or abnormally terminates before issuing CLOSE, the CLOSE macro instruction will be issued by the control program, with control-program information in these registers.

The parameter list pointed to by register 1 is a 16-byte area aligned on a fullword boundary. Figure 16 shows the contents of the area.
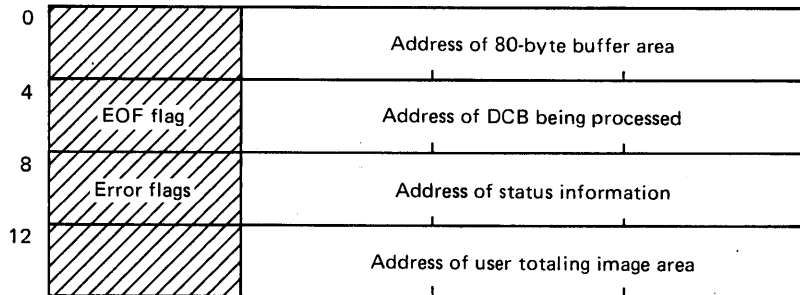


Figure 16. Parameter List Passed to User Label Exit Routine

The first address in the parameter list points to an 80-byte label buffer area. For input, the control program reads a user label into this area before passing control to the label routine. For output, the user label exit routine constructs labels in this area and returns to the control program, which writes the label. When an input trailer label routine receives control, the EOF flag (high-order byte of the second entry in the parameter list) is set as follows:

    bit 0 = 0: Entered at end-of-volume
    bit 0 = 1: Entered at end-of-file
    bits 1-7: Reserved

When a user label exit routine receives control after an uncorrectable I/O error has occurred, the third entry of the parameter list contains the address of the standard status information. The error flag (high-order byte of the third entry in the parameter list) is set as follows:

    bit 0 = 1: Uncorrectable I/O error
    bit 1=1: Error occurred during writing of updated label
    bits 2-7: Reserved

The fourth entry in the parameter list is the address of the user totaling image area. This image area is the entry in the user totaling save area that corresponds to the last record physically written on the volume. The image area is discussed further under "User Totaling."

Each routine must create or verify one label of a header or trailer label group, place a return code in register 15, and return control to the operating system. The operating system responds to the decimal return code as shown in Figure 17.

You can create user labels only for data sets on magnetic-tape volumes with IBM standard labels or American National Standard labels and for data sets on direct-access volumes. When you specify both user labels and IBM standard or American National Standard labels in a DD statement by specifying LABEL=(,SUL) or LABEL=(,AUL)

| Routine Type | Return Code | System Response |
|---|---|---|
| Input header or trailer label | 0 | Normal processing is resumed. If there are any remaining labels in the label group, they are ignored. |
| | 4 | The next user label is read into the label buffer area and control is returned to the exit routine. If there are no more labels in the label group, normal processing is resumed. |
| | 8[1] | The label is written from the label buffer area and normal processing is resumed. |
| | 12[1] | The label is written from the label area, the next label is read into the label buffer area, and control is returned to the label processing routine. If there are no more labels, processing is resumed. |
| Output header or trailer label | 0 | Normal processing is resumed; no label is written from the label buffer area. |
| | 4 | User label is written from the label buffer area. Normal processing is resumed. |
| | 8 | User label is written from the label buffer area. If fewer than eight labels have been created, control is returned to the exit routine, which then creates the next label. If eight labels have been created, normal processing is resumed. |

[1] Your input label routines can only return these codes when you are processing a physical sequential data set opened for UPDAT or a direct data set opened for OUTPUT or UPDAT. These return codes allow.you to verify the existing labels, update them if necessary, then request that the system write the updated labels.

Figure 17. System Response to a User Label Exit Routine Return Code

and there is an active entry in the exit list, a label exit is always taken. Thus, a label exit is taken even when an input data set does not contain user labels, or when no user label track has been allocated for writing labels on a direct-access volume. In either case, the appropriate exit routine is entered with the buffer area address parameter set to 0. On return from the exit routine, normal processing is resumed; no return code is necessary.

Label exits are not taken for system output (SYSOUT) data sets, or for data sets on volumes that do not have standard labels. For other data sets, exits are taken as follows:

- When an input data set is opened, the input header label exit 01 is taken. If the data set is on tape being opened for RDBACK, user trailer labels will be processed.

- When an output data set is opened, the output header label exit 02 is taken. However, if the data set already exists and DISP=MOD is coded in the DD statement, the input trailer label exit 03 is taken to process any existing trailer labels. If the input trailer label exit 03 does not exist, then the deferred input trailer label exit 0C is taken if it exists; otherwise, no label exit is taken. For tape, these trailer labels will be overwritten by the new output data or by EOV or close processing when writing new standard trailer labels. For direct-access devices, these trailer labels will still exist unless rewritten by EOV or close processing in an output trailer label exit.

- When an input data set reaches end-of-volume, the input trailer label exit 03 is taken. If the data set is on tape opened for RDBACK, header labels will be processed. The input trailer label exit 03 is not taken if you issue an FEOV macro instruction. If a defer input trailer label exit 0C is present, it is taken instead of any input trailer label exit 03. After switching volumes, the input header label exit 01 is taken. If the data set is on tape opened for RDBACK, trailer labels will be processed.

- When an output data set reaches end-of-volume, the output trailer label exit 04 is taken. After switching volumes, output header label exit 02 is taken.

- When an input data set reaches end-of-data, the input trailer label exit 03 is taken before the EODAD exit, unless the DCB exit list contains a defer input trailer label exit 0C.

- When an input data set is closed, no exit is taken unless the data set was previously read to end-of-data and the defer input trailer label exit 0C is present. If so, the defer input trailer label exit 0C is taken to process trailer labels, or if the tape is opened for RDBACK, header labels.

- When an output data set is closed, the output trailer label exit 04 is taken.

To process records in reverse order, a data set on magnetic tape can be read backward. When you read backward, header label exits are taken to process trailer labels, and trailer label exits are taken to process header labels. The system presents labels from a label group in ascending order by label number, which is the order in which the labels were created. If necessary, an exit routine can determine label type (UHL or UTL) and number by examining the first four characters of each label. Tapes with IBM standard labels and direct-access devices can have as many as eight user labels. Tapes with American National Standard labels can have unlimited user labels.

If an uncorrectable error occurs during reading or writing of a user label, the system passes control to the appropriate exit routine with the third word of the parameter list flagged and pointing to status information.

After an input error, the exit routine must return control with an appropriate return code (0 or 4). No return code is required after an output error. If an output error occurs while the system is opening a data set, the data set is not opened (DCB is flagged) and control is returned to your program. If an output error occurs at any other time, the system attempts to resume normal processing.

*User Totaling (BSAM and QSAM only):* When creating or processing a data set with user labels, you may develop control totals for each volume of the data set and store this information in your user labels. For example, a control total that was accumulated as the data set was created can be stored in your user label and later compared with a total accumulated during processing of the volume. User totaling assists you by synchronizing the control data you create with records physically written on a volume. For an output data set without user labels, you can also develop a control total that will be available to your end-of-volume routine.

To request user totaling, you must specify OPTCD=T in the DCB macro instruction or in the DCB parameter of the DD statement. The area in which you accumulate the control data (the user totaling area) must be identified to the control program by an entry of hexadecimal 0A in the DCB exit list. OPTCD=T cannot be specified for SYSIN or SYSOUT data sets.

The user totaling area, an area in storage that you provide, must begin on a halfword boundary and be large enough to contain your accumulated data plus a 2-byte length field. The length field must be the first 2 bytes of the area and specify the length of the entire area. A data set for which you have specified user totaling (OPTCD=T) will not be opened if either the totaling area length or the address in the exit list is 0, or if there is no X'0A' entry in the exit list.

The control program establishes a user totaling save area, in which the control program preserves an image of your totaling area, when an I/O operation is scheduled. When the output user label exits are taken, the address of the save area entry (user totaling image area) corresponding to the last record physically written on a volume is passed to you in the fourth entry of the user label parameter list. This parameter list is described in the section "Standard User Label Exit." When an end-of-volume exit is taken for an output data set and user totaling has been specified, the address of the user totaling image area is in register 0.

When using user totaling for an output data set, that is, when creating the data set, you must update your control data in your totaling area before issuing a PUT or a WRITE macro instruction. The control program places an image of your totaling area in the user

totaling save area when an I/O operation is scheduled. A pointer to the save area entry (user totaling image area) corresponding to the last record physically written on the volume, is passed to you in your label processing routine. Thus you can include the control total in your user labels. When subsequently using this data set for input, you can accumulate the same information as you read each record and compare this total with the one previously stored in the user trailer label. If you have stored the total from the preceding volume in the user header label of the current volume, you can process each volume of a multivolume data set independently and still maintain this system of control.

When variable-length records are specified with the totaling facility for user labels, special considerations are necessary. Since the control program determines whether a variable-length record will fit in a buffer after a PUT or a WRITE has been issued, the total you have accumulated may include one more record than is actually written on the volume. In the case of variable-length spanned records, the accumulated total will include the control data from the volume-spanning record although only a segment of the record is on that volume. However, when you process such a data set, the volume-spanning record or the first record on the next volume will not be available to you until after the volume switch and user label processing are completed. Thus the totaling information in the user label may not agree with that developed during processing of the volume.

One way you can resolve this situation is to maintain, when you are creating a data set, control data pertaining to each of the last two records and include both totals in your user labels. Then the total related to the last complete record on the volume and the volume-spanning record or the first record on the next volume would be available to your user label routines. During subsequent processing of the data set, your user label routines can determine if there is agreement between the generated information and one of the two totals previously saved.

*Data Control Block Open Exit:* You can specify in an exit list the address of a routine that completes or modifies a DCB and does any additional processing required before the data set is completely open. The routine is entered during the opening process after the JFCB has been used to supply information for the DCB. The routine can determine data set characteristics by examining fields completed from the data set labels. When your DCB exit routine receives control, the three, low-order bytes of register 1 will contain the address of the DCB currently being processed.

As with label processing routines, register 14's contents must be preserved and restored if any macro instructions are used in the routine. Control is returned to the operating system by a RETURN macro instruction; no return code is required.

This exit is mutually exclusive with the JFCBE exit. If you need both the JFCBE and data control block exits, you must use the JFCBE exit to pass control to your routines.

*QSAM Parallel Input Exit:* A request for parallel input processing is indicated by including the address of a parallel data access block (PDAB) in the DCB exit list. The address must be on a fullword boundary with the first byte of the entry containing X'12' or, if it is the last entry, X'92'. For more information on parallel input processing, see "Parallel Input Processing (QSAM Only)".

*JFCBE Exit:* JCL-specified setup requirements for the 3800 printer cause a JFCB extension (JFCBE) to be created to reflect those specifications. The JFCBE exit can be used to examine or modify those specifications in the JFCBE. You can provide a JFCBE exit routine to examine or modify those specifications. The address of the routine should be placed in an exit list. This exit is taken during open processing and is mutually exclusive with the data control block exit. If you need both the JFCBE and data control block exits, you must use the JFCBE exit to pass control to your routines.

When control is passed to your exit routine, the contents of register 0 and 1 will be:

**Register   Contents**

0        If a JFCBE exists, this register will point to an area in your storage into which a copy of the JFCBE has been placed. If a JFCBE does not exist, this register will be zero.

1        The address of the DCB being processed.

Registers 2-15 will contain the standard user exit contents.

The area in your storage pointed to by register 0 will also contain the 4-byte FCB identification which is obtained from the JFCB. The FCB identification is placed in the four bytes following the 176-byte JFCBE.

If your copy of the JFCBE is modified during an exit routine, you should indicate this fact by turning on bit JFCBEOPN (X'80' in JFCBFLAG) in the JFCBE copy. On return to open, this bit indicates whether the system copy is to be updated. The 4-byte FCB identification in your area will be used to update the JFCB regardless of the bit setting. Checkpoint/restart also interrogates this bit to determine which version of the JFCBE will be used at restart time. If this bit is not on, the JFCBE generated by the restart JCL will be used.

***End-of-Volume Exit:*** You can specify in an exit list the address of a routine that is entered when end-of-volume is reached in processing of a physical sequential data set.

When the end-of-volume routine is entered, register 0 contains 0 unless user totaling was specified. If you specified user totaling in the DCB macro instruction (by coding OPTCD=T) or in the DD statement for an output data set, register 0 contains the address of the user totaling image area. The routine is entered after a new volume has been mounted and all necessary label processing has been completed. If the volume is a reel of magnetic tape, the tape is positioned after the tapemark that precedes the beginning of the data.

You can use the end-of-volume (EOV) exit routine to take a checkpoint by issuing the CHKPT macro instruction, which is discussed in *OS/VS Checkpoint/Restart*; specifications for the CHKPT macro are also included in *OS/VS1 Data Management Macro Instructions*. If a checkpointed job step terminates abnormally, it can be restarted from the EOV checkpoint. When the job step is restarted, the volume is mounted and positioned as upon entry to the routine. Restart becomes impossible if changes are made to the supervisor call (SVC) library between the time the checkpoint is taken and the time the job step is restarted. When the step is restarted, pointers to end-of-volume modules must be the same as when the checkpoint was taken.

The end-of-volume exit routine returns control in the same manner as the data control block exit routine. Register 14's contents must be preserved and restored if any macro instructions are used in the routine. Control is returned to the operating system by a RETURN macro instruction; no return code is required.

***Block Count Exit:*** You can specify in an exit list the address of a routine that will allow you to abnormally terminate the task or continue processing when the end-of-volume routine finds an unequal block count condition. When you are using standard labeled input tapes, the block count in the trailer label is compared by the end-of-volume routine with the block count in the DCB. The count in the trailer label reflects the number of blocks written when the data set was created. The number of blocks read when the tape is used as input is contained in the DCBBLKCT field of the DCB.

The routine is entered during end-of-volume processing. The trailer label block count is passed in register 0. You may gain access to the count field in the DCB by using the address passed in register 1 plus the proper displacement, as given in *OS/VS1 System Data Areas*. If the block count in the DCB differs from that in the trailer label when no

exit routine is provided, the task is abnormally terminated. The routine must terminate with a RETURN macro instruction and a return code that indicates what action is to be taken by the operating system, as shown in Figure 18. As with other exit routines, register 14's contents must be saved and restored if any macro instructions are used.

| Return Code | System Action |
| --- | --- |
| 0 | The task is to be abnormally terminated. |
| 4 | Normal processing is to be resumed. |

Figure 18. System Response to Block Count Exit Return Code

*Defer Nonstandard Input Trailer Label Exit:* In an exit list, you can specify a code that indicates that you want to defer nonstandard input trailer label processing from end-of-data until the data set is closed. The address portion of the entry is not used by the operating system.

An end-of-volume condition exists in several situations. Two examples are: (1) when the system reads a filemark or tapemark at the end of a volume of a multivolume data set but that volume is not the last, and (2) when the system reads a filemark or tapemark at the end of a data set. The first situation is referred to here as an end-of-volume condition, and the second as an end-of-data condition, although it, too, can occur at the end of a volume.

For an end-of-volume (EOV) condition, the EOV routine passes control to your nonstandard input trailer label routine, whether or not this exit code is specified. For an end-of-data condition when this exit code is specified, the EOV routine does not pass control to your nonstandard input trailer label routine. Instead, the close routine passes control to your end-of-data routine.

*FCB Image Exit:* You can specify in an exit list the address of a forms control buffer (FCB) image. This FCB image can be loaded into the forms control buffer of the printer control unit. The FCB controls the movement of forms in printers that do not use a carriage control tape.

Multiple exit list entries in the exit list can define FCBs. The open and SETPRT routines search the exit list for requested FCBs before searching SYS1.IMAGELIB.

The first 4 bytes of the FCB image contain the image identifier. To load the FCB, this image identifier is specified in the FCB parameter of the DD statement, by the SETPRT macro instruction, or by the system operator in response to message IEC127D or IEC129D.

The image identifier is followed by the FCB load module, described in *OS/VS1 Data Management for System Programmers*.

You can use an exit list to define an FCB image only when writing to an online printer. Figure 19 illustrates one way the exit list can be used to define an FCB image.

*DCB ABEND Exit:* The DCB ABEND exit is provided to give you some options regarding the action you want the system to take when a condition arises that may result in abnormal termination of your task. This exit can be taken any time an ABEND condition arises during the process of opening, closing, or handling an end-of-volume condition for a DCB associated with your task.

When an ABEND condition arises, a write-to-programmer message about the ABEND is issued and your DCB ABEND exit is given control, provided there is an active DCB ABEND exit routine address in the DCB being processed. If STOW called the end-of-volume routines to get secondary space to write an end-of-file mark for a partitioned data set, the DCB ABEND exit routine will not be given control if an ABEND condition occurs. The contents of the registers when your exit routine is entered

```
                ...
                DCB       ..,EXLST=EXLIST
                ...
EXLIST          DS        0F
                DC        X'10'              Flag code for FCB image
                DC        AL3(FCBIMG)        Address of FCB image
                DC        X'80000000'        End of EXLST and a null entry
FCBIMG          DC        CL4'IMG1'          FCB identifier
                DC        X'00'              FCB is not a default
                DC        AL1(67)            Length of FCB
                DC        X'90'              Offset print line
* 16 line character positions to the right
                DC        X'00'              Spacing is 6 lines per inch
                DC        5X'00'             Lines 2-6 no channel codes
                DC        X'01'              Line 7 channel 1
                DC        6X'00'             Lines 8-13 no channel codes
                DC        X'02'              Line (or Lines) 14 channel 2
                DC        5X'00'             Line (or Lines) 15-19 no channel codes
                DC        X'03'              Line (or Lines) 20 channel 3
                DC        9X'00'             Line (or Lines) 21-29 no channel codes
                DC        X'04'              Line (or Lines) 30 channel 4
                DC        19X'00'            Line (or Lines) 31-49 no channel codes
                DC        X'05'              Line (or Lines) 50 channel 5
                DC        X'06'              Line (or Lines) 51 channel 6
                DC        X'07'              Line (or Lines) 52 channel 7
                DC        X'08'              Line (or Lines) 53 channel 8
                DC        X'09'              Line (or Lines) 54 channel 9
                DC        X'0A'              Line (or Lines) 55 channel 10
                DC        X'0B'              Line (or Lines) 56 channel 11
                DC        X'0C'              Line (or Lines) 57 channel 12
                DC        8X'00'             Line (or Lines) 58-65 no channel codes
                DC        X'10'              End of FCB image
                ...
                END
//ddname        DD        UNIT=3211, FCB=(IMG1,VERIFY)
/*
```

Figure 19. Defining an FCB Image for a 3211

are the same as for other DCB exit list routines except that the three, low-order bytes of register 1 contain the address of the parameter list described in Figure 20. Your ABEND exit routine can choose one of four options:

- to immediately terminate your task,

- to delay the ABEND until all of the DCBs in the same OPEN or CLOSE macro instruction are opened or closed,

- to ignore the ABEND condition and continue processing without making reference to the DCB on which the ABEND condition was encountered, or

- to try to recover from the error.

Not all of these options are available for each ABEND condition. Your DCB ABEND exit routine must determine which option is available by examining the contents of the option mask byte (byte 3) of the parameter list. The address of the parameter list is passed in register 1. Figure 20 shows the contents of the parameter list and the possible settings of the option mask when your routine receives control. All information in the parameter list is in binary.

| Bit | Meaning |
|-----|---------|
| 0-3 | Reserved for Future Use |
| 4 | OK to Recover |
| 5 | OK to Ignore |
| 6 | OK to Delay |
| 7 | Reserved for Future Use |

**Fullword Boundary**

Displacement

| | | | |
|---|---|---|---|
| 0 | System Completion Code* | Return Code | Option Mask |
| 4 | DCB Address | | |
| 8 | Open/Close/End-of-Volume Work Area Address | | |
| 12 | 00 | Recovery Work Area Address | |

*In the first 12 bits.

Figure 20. Parameter List Passed to DCB ABEND Exit Routine

When your DCB ABEND exit routine returns control to the system control program (this can be done using the RETURN macro instruction), the option mask byte should contain the setting that specifies the action you want to take. These actions and the corresponding settings of the option mask byte are:

**Bit Setting   Action**

0       abnormally terminate the task immediately

4       ignore the ABEND condition

8       delay the ABEND until the other DCBs being processed concurrently are opened or closed

12      make an attempt to recover

You must inspect bits 4, 5, and 6 of the option mask byte (byte 3 of the parameter list) to determine which options are available. If a bit is set to 1, the corresponding option is available. Indicate your choice by inserting the appropriate value in byte 3 of the parameter list, overlaying the bits you inspected. If you use a value that specifies an option that is not available, the ABEND is issued immediately.

If the contents of the option mask are 0, you must request an immediate ABEND by leaving the value of 0 in the option mask unchanged.

If bit 5 of the option mask is set to 1, you can ignore the ABEND by placing a decimal value of 4 in byte 3 of the parameter list. Processing on the current DCB stops. If you subsequently attempt to use this DCB, the results are unpredictable. If you ignore an

error in end-of-volume, the data set will be closed before control is returned to your program at the point which caused the end-of-volume condition (unless the end-of-volume routines were called by the close routines). If the end-of-volume routines were called by the close routines, an ABEND macro will be issued even though the ignore option was selected.

If bit 6 of the option mask is set to 1, you can delay the ABEND by placing a decimal value of 8 in byte 3 of the parameter list. All other DCBs waiting for open or close processing will be processed before the ABEND is issued. For end-of-volume, however, you can't delay the ABEND because the end-of-volume routine never has more than one DCB to process.

If bit 4 of the option mask is set to 1, you can attempt to recover. Place a decimal value of 12 in byte 3 of the parameter list and provide information for the recovery attempt. Figure 21 lists the ABEND conditions for which recovery can be attempted. For ABEND conditions which can be ignored or delayed, see *OS/VS Message Library: VS1 System Messages*.

| System Completion Code | Return Code | Description of Error |
|---|---|---|
| 213 | 04 | DSCB was not found on volume specified. |
| 237 | 04 | Block count in DCB does not agree with block count in trailer label. |
| 413 | 18 | Data set was opened for input and no volume serial number was specified. |
| 613 | 08 | I/O error occurred during reading of tape label. |
| | 0C | Invalid tape label was read. |
| | 10 | I/O error occurred during writing of tape label. |
| | 14 | I/O error occurred during writing of tapemark following header labels. |
| 713 | 04 | A data set on magnetic tape was opened for INOUT, but the volume contained a data set whose expiration date had not been reached and the operator denied permission. |
| 717 | 10 | I/O error occurred during reading of trailer label 1 to update block count in DCB. |
| 813 | 04 | Data set name on header label does not match data set name on DD statement. |

Figure 21. Conditions for which Recovery Can Be Attempted

*Recovery Requirements:* For the recovery attempt, you should supply a recovery work area (see Figure 22) with a new volume serial number for each volume associated with an error. If no new volumes are supplied, recovery will be attempted with the existing volumes, but the likelihood of successful recovery is greatly reduced.

If you request recovery for system completion code 213, return code 04, you must indicate in your job control language (JCL) that the volumes are nonsharable by specifying unit affinity, deferred mounting, or more volumes than units for the data set.

If you request recovery for system completion code 237, return code 04, you don't need to supply new volumes or a work area. The condition that caused the ABEND is the disagreement between the block count in the DCB and that in the trailer label. This disagreement is ignored to permit recovery.

If you request recovery for system completion code 717, return code 10, you don't need to supply new volumes or a work area. The ABEND is caused by an I/O error during updating of the DCB block count. To permit recovery, the block count is not updated. Consequently, an abnormal termination with system completion code 237, return code 04, may result when you try to read from the tape after recovery. You may attempt

| Bit | Meaning |
|-----|---------|
| 0 | Free This Work Area |
| 1 | Volume Serial Numbers Are Provided |
| 2-7 | Reserved for Future Use |

Halfword Boundary

Displacement

| | | | |
|---|---|---|---|
| 0 | Length of This Work Area | Option Byte | Subpool Number |
| 4 | Number of New Volumes | New Volume Serial Numbers (6 bytes each) | |
| 8 | | | |

Figure 22. Recovery Work Area

recovery from the ABEND with system completion code 237, return code 04, as explained in the preceding paragraph.

System completion codes and their associated return codes are described in *OS/VS Message Library: VS1 System Codes.*

The work area that you supply for the recovery attempt must begin on a halfword boundary and can contain the information described in Figure 22. Place a pointer to the work area in the last 3 bytes of the parameter list pointed to by register 1 and described in Figure 20.

If you acquire the storage for the work area by using the GETMAIN macro instruction, you can request that it be freed by a FREEMAIN macro instruction after all information has been extracted from it. Set the high-order bit of the option byte in the work area to 1 and place the number of the subpool from which the work area was requested in byte 3 of the recovery work area.

Only one recovery attempt per data set is allowed during open, close, or end-of-volume processing. If a recovery attempt is unsuccessful, you may not request another recovery. The second time through the exit routine you may request only one of the other options (if allowed): issue the ABEND immediately, ignore the ABEND, or delay the ABEND. If at any time you select an option that is not allowed, the ABEND is issued immediately.

Note that if recovery is successful, you still receive an ABEND message on your listing. This message refers to the ABEND that would have been issued if the recovery had not been successful.

## Modifying the Data Control Block

You can complete or modify the DCB during execution of your program. You can also determine data set characteristics from information supplied by the data set labels. Changes or additions can be made before opening of the data set, after closing it, during the DCB exit routine, or while the data set is open. Naturally, any information must be supplied before it is needed.

Because each DCB does not have a symbolic name for each field, a DCBD macro instruction must be used to supply the symbolic names. By loading a base register with the address of the DCB to be processed, you can refer to any field symbolically.

The DCBD macro instruction generates a dummy control section (DSECT) named IHADCB. The name of each field consists of DCB followed by the first five letters of the keyword operand that represents the field in the DCB macro instruction. For example, the field reserved for blocksize is referred to as DCBBLKSI. For the names of other fields, including names of bits, see *OS/VS1 Data Management Macro Instructions.*

The attributes of each DCB field are defined in the dummy control section. Because each field in the DCB is not necessarily aligned on a fullword boundary, care must be taken when storing or moving data into the field. The length attribute and the alignment of each field can be determined from an assembly listing of the DCBD macro instruction.

The DCBD macro instruction can be coded once to describe all DCBs even though their fields differ because of differences in data set organization and access technique. It must not be coded more than once for a single assembly. If it is coded before the end of a control section, it must be followed by a CSECT or DSECT statement to resume the original control section.

**Changing an Address in the Data Control Block:** Figure 23 illustrates how you can modify a field in the data control block. The DCBD macro instruction defines the symbolic name of each field.

```
              . . .
              OPEN    ( TEXTDCB , INOUT )
              . . .
EOFEXIT       CLOSE       ( TEXTDCB , REREAD ) , TYPE=T
              LA          10 , TEXTDCB
              USING       IHADCB , 10
              MVC         DCBSYNAD+1 ( 3 ) , =AL3 ( OUTERROR )
              B           OUTPUT
INERROR       STM         14 , 12 , SYNADSA+12
              . . .
OUTERROR      STM         14 , 12 , SYNADSA+12
              . . .
TEXTDCB       DCB         DSORG=PS , MACRF=( R , W ) , DDNAME=TEXTTAPE ,        C
                      EODAD=EOFEXIT , SYNAD=INERROR
              DCBD        DSORG=PS
              . . .
```

Figure 23. Modifying a Field in the Data Control Block

The data set defined by the data control block TEXTDCB is opened for use as both an input and an output data set. When its use as an input data set is completed, the EODAD routine closes the data set temporarily to reposition the volume for output. The EODAD routine then uses the dummy control section IHADCB to change the error exit address (SYNAD) from INERROR to OUTERROR.

The EODAD routine loads the address TEXTDCB into register 10, which it uses as a base register for IHADCB. It then moves the address OUTERROR into the DCBSYNAD field of the DCB. This field is a fullword, but contains information that

must not be disturbed in the high-order byte. For this reason, care must be taken to change only the 3 low-order bytes of the field.

All unused address fields in the DCB except DCBEXLST are set to 1 during the DCB macro expansion. Many system routines interpret a value of 1 in an address field to mean "no address specified." If you modify an address field and then want to reset it to "no address specified," you should set it to a value of 1.

## Sharing a Data Set

There are two conditions under which a data set on a direct-access device can be shared by two or more tasks:

- Two or more DCBs are opened and used concurrently by the tasks to refer to the same, shared data set (multiple DCBs).

- Only one DCB is opened and used concurrently by multiple tasks in a single job step (a single, shared DCB).

Job control language (JCL) statements and macro instructions are provided in the operating system to help you to ensure the integrity of the data sets you wish to share among the tasks that process them. Figures 24 and 25 show which JCL and macro instructions you should use, depending on the access method your task is using and mode of access (input, output, or update).

Figure 24 describes the macro instructions, JCL, and processing procedures you should use if more than one DCB has been opened to the shared data set. The DCBs can be used by tasks in the same or different job steps.

MULTIPLE DCBs

| Access Mode | Access Method | | | | |
|---|---|---|---|---|---|
| | BSAM, BPAM, BDAM Create | QSAM | BDAM | QISAM | BISAM |
| Input | DISP = SHR | DISP = SHR | DISP = SHR | DISP = SHR | DISP = SHR |
| Output | No Facility | No Facility | DISP = SHR | No Facility | DISP = SHR and ENQ on Data Set |
| Update | DISP = SHR and ENQ on Block | DISP = SHR and Guarantee Discrete Blocks | DISP = SHR and ENQ on Block | DISP = SHR and ENQ on Data Set and Guarantee Discrete Blocks | DISP = SHR and ENQ on Data Set and Guarantee Discrete Blocks |

**DISP=SHR:**

Each job step sharing an existing data set must code SHR as the subparameter of the DISP parameter on the DD statement for the shared data set to allow the steps to execute concurrently. For additional information about ensuring data set integrity, see *OS/VS1 JCL Services.* If the tasks are in the same job step, DISP=SHR is not required.

**No Facility:**

There are no facilities in the operating system for sharing a data set under these conditions.

**ENQ on Data Set:**

In addition to coding DISP=SHR on the DD statement for the data set that is to be shared, each task must issue ENQ and DEQ macro instructions naming the data set as resource for which exclusive control is required. The ENQ must be issued before the GET (READ); the DEQ macro should be issued after the PUTX or CHECK macro that concludes the operation. See *OS/VS1 Supervisor Services and Macro Instructions* for additional information on the use of ENG and DEQ macro instructions.

**Guarantee Discrete Blocks:**

When you are using the access methods that provide blocking and unblocking of records (QSAM, QISAM, and BISAM), it is necessary that every task updating the data set ensure that it is not updating a block that contains a record being updated by any other task. There are no facilities in the operating system for ensuring that discrete blocks are being processed by different tasks.

**ENQ on Block:**

If you are updating a shared data set (specified by coding DISP=SHR on the DD statement) using BSAM or BPAM, your task and all other tasks must serialize processing of each block of records by issuing an ENQ macro instruction before the READ macro and a DEQ macro after the CHECK macro that follows the WRITE macro you issued to update the record. If you are using BDAM, the same procedure may be used; however BDAM provides for enqueuing on a block of records using the READ exclusive option, which is requested by coding MACRF=X in the DCB and an X in the type operand of the READ and WRITE macro instructions. See "Exclusive Control for Updating" for an example of the use of the BDAM macros.

Figure 24. JCL, Macro Instructions, and Procedures Required to Share a Data Set Using Multiple DCBs.

Figure 25 describes the macros you can use to serialize processing of a shared data set when a single DCB is being shared by several tasks in a job step. The DISP=SHR specification on the DD statement is not required.

Data sets can also be shared both ways at the same time: more than one DCB can be opened for a shared data set, while more than one task can be'sharing one of the DCBs. Under this condition, the serialization techniques specified for indexed sequential and direct data sets in the Figure 24 satisfy the requirement. For sequential and partitioned data sets, the techniques specified in Figure 24 and Figure 25 must be used.

**A SINGLE SHARED DCB**

| Access Mode | Access Method | | | | |
|---|---|---|---|---|---|
| | BSAM, BPAM, BDAM Create | QSAM | BDAM | QISAM | BISAM |
| Input | ENQ | ENQ | No Action Required | ENQ | ENQ |
| Output | ENQ | ENQ | No Action Required | ENQ and Key Sequence | ENQ |
| Update | ENQ | ENQ | ENQ on Block | ENQ | ENQ |

**ENQ:**
When a data set is being shared by two or more tasks in the same job step (all of which must be using the same DCB), each task processing the data set must issue an ENQ macro instruction on a predefined resource name before issuing the macro or macros that begin the input/output operation. Each task must also release exclusive control by issuing the DEQ macro instruction at the next sequential instruction following the input/output macro. If, however, you are processing an indexed sequential data set sequentially using the SETL and ESETL macros, you must issue the ENQ macro before the SETL macro and the DEQ macro after the ESETL macro. Note also that if two tasks are writing different members of a partitioned data set, each task should issue the ENQ macro instruction before the FIND macro and issue the DEQ macro after the STOW macro that completes processing of the member. Additional reference information on the ENQ and DEQ macros is presented in *OS/VS1 Supervisor Services and Macro Instructions*. For an example of the use of ENQ and DEQ macro instructions with BISAM, see Figure 59.

**No Action Required:**
Sharing a Direct Data Set: BDAM supports multiple task users of a single DCB when working with existing data sets. When operating in load mode, however, only one task may use the DCB at a time. The following restrictions and comments apply when operating in a multitasking mode with existing data sets:

- Subpool 0 must be shared.

- The user should insure that a WAIT or CHECK macro has been issued for all outstanding BDAM requests before the task issuing the READ or WRITE macro terminates. In case of abnormal termination this can be done through a STAE/STAI or ESTAE exit.

- FREEDBUF and/or RELEX macros should be issued to free any resources that could still be held by the terminating task. This can be done during or after task termination.

**ENQ on Block:**
When updating a shared BDAM data set, every task must use the BDAM exclusive control option, which is requested by coding MACRF=X in the DCB macro and an X in the type operand of the READ and WRITE macro instructions. See "Exclusive Control for Updating" in this book for an example of the use of BDAM macros. Note that all tasks sharing a data set must share subpool 0 (see the ATTACH macro description in *OS/VS1 Supervisor Services and Macro Instructions*.

**Key Sequence:**
Tasks sharing a QISAM load-mode DCB must ensure that the records to be written are presented in ascending key sequence; otherwise, a sequence check will result in (1) control being passed to the SYNAD routine identified by the DCB, or (2) if there is no SYNAD routine, termination of the task.

Figure 25. Macro Instructions and Procedures Required to Share a Data Set Using a Single DCB

More information on opening and closing data sets by more than one task is contained in Part 2, "Opening and Closing a Data Set."

**Shared Direct-Access Storage Devices:** At some installations, a direct-access storage device is shared by two or more independent computing systems. Tasks executed on these systems can share data sets stored on the device. For details, refer to *OS/VS1 Planning and Use Guide.*

# PART 2: DATA MANAGEMENT PROCESSING PROCEDURES

## Data-Processing Techniques

The operating system allows you to concentrate most of your efforts on processing the records read or written by the data management routines. To get the records read and written, your main responsibilities are to describe the data set to be processed, the buffering techniques to be used, and the access method. An access method has been defined as the combination of data set organization and the technique used to gain access to the data. Data access techniques are discussed here in two categories—queued and basic.

### Queued Access Technique

The queued access technique provides GET and PUT macro instructions for transmitting data within virtual storage. These macro instructions cause automatic blocking and deblocking of the records stored and retrieved. Anticipatory (look-ahead) buffering and synchronization (overlap) of input and output operations with central processing unit (CPU) processing are automatic features of the queued access technique.

Because the operating system controls buffer processing, you can use as many input/output (I/O) buffers as needed without reissuing GET or PUT macro instructions to fill or empty buffers. Usually, more than one input block is in storage at any given time, so I/O operations do not delay record processing.

Because the operating system synchronizes input/output with processing, you need not test for completion, errors, or exceptional conditions. After a GET or PUT macro instruction is issued, control is not returned to your program until an input area is filled or an output area is available. Exits to error analysis (SYNAD) and end-of-volume or end-of-data (EODAD) routines are automatically taken when necessary.

#### GET—Retrieve a Record

The GET macro instruction obtains a record from an input data set. It operates in a logical sequential and device-independent manner. As required, the GET macro instruction schedules the filling of input buffers, deblocks records, and directs input error recovery procedures. For sequential data sets, it also merges record segments into logical records. After all records have been processed and the GET macro instruction detects an end-of-data indication, the system automatically checks labels on sequential data sets and passes control to your end-of-data (EODAD) routine. If an end-of-volume condition is detected for a sequential data set, the system provides automatic volume switching if the data set extends across several volumes or if concatenated data sets are being processed. If you specify OPTCD=Q in the DCB, GET causes input data to be translated from ASCII to EBCDIC.

#### PUT—Write a Record

The PUT macro instruction places a record into an output data set. Like the GET macro instruction, it operates in a logical sequential and device-independent manner. As required, the PUT macro instruction schedules the emptying of output buffers, blocks records, and handles output error correction procedures. For sequential data sets, it also initiates automatic volume switching and label creation, and also segments records for spanning. If you specify OPTCD=Q in the DCB, PUT causes output to be translated from EBCDIC to ASCII.

If the PUT macro instruction is directed to a card punch or printer, the system automatically adjusts the number of records or record segments per block of format-F or format-V blocks to 1. Thus, you can specify a record length (LRECL) and blocksize (BLKSIZE) to provide an optimum blocksize if the records are temporarily placed on magnetic tape or a direct-access volume.

For spanned variable-length records, the blocksize must be equivalent to the length of one card or one print line. Record size may be greater than blocksize in this case.

## PUTX—Write an Updated Record

The PUTX macro instruction is used to update a data set or to create an output data set using records from an input data set as a base. PUTX updates, replaces, or inserts records from existing data sets but does not create records.

When you use the PUTX macro instruction to update, each record is returned to the data set referred to by a previous locate mode GET macro instruction. The buffer containing the updated record is flagged and written back to the same location on the direct-access storage device from which it was read. The block is not written until a GET macro instruction is issued for the next buffer, except when a spanned record is to be updated. In that case, the block is written with the next GET macro instruction.

When the PUTX macro instruction is used to create an output data set, you can add new records by using the PUT macro instruction. As required, the PUTX macro instruction blocks records, schedules the writing of output buffers, and handles output error correction procedures.

## Parallel Input Processing (QSAM Only)

QSAM parallel input processing may be used to process two or more input data sets concurrently, such as sorting or merging several data sets at the same time. This eliminates the need for issuing a separate GET macro instruction to each DCB processed. The get routine for parallel input processing selects a DCB with a ready record and then transfers control to the normal get routine. If there is no DCB with a ready record, a multiple WAIT macro instruction is issued.

Parallel input processing provides a logical input record from a queue of data sets with equal priority. The function supports QSAM with input processing, simple buffering, locate or move mode, and fixed, variable, or undefined length records. Spanned records, track-overflow records, dummy data sets, and SYSIN data sets are not supported.

Parallel input processing can be interrupted at any time to retrieve records from a specific data set, or to issue control instructions to a specific data set. When the retrieval process has been completed, parallel input processing may be resumed.

Data sets can be added to or deleted from the data set queue at any time. It is important to note, however, that as each data set reaches an end-of-data condition, the data set must be removed from the queue with the CLOSE macro instruction before a subsequent GET macro instruction is issued for the queue; otherwise, the task may be terminated abnormally.

A request for parallel input processing is indicated by including the address of a parallel data access block (PDAB) in the DCB exit list. For additional information on the DCB exit list, see "Exit List (EXLST)."

With the use of the PDAB macro instruction, you can create and format a work area that identifies the maximum number of DCBs that can be processed at any one time. If you exceed the maximum number of entries indicated in the PDAB macro when adding a DCB to the queue with the OPEN macro, the data set will not be available for parallel input processing; however, it may be available for sequential processing.

When issuing a parallel GET macro, register 1 must always point to a PDAB. You may load the register or let the GET macro do it for you. When control is returned to you, register 1 contains the address of a logical record from one of the data sets in the queue; registers 2-13 contain their original contents at the time the GET macro was issued; registers 14, 15, and 0 are changed. You can locate the data set from which the record was retrieved through the PDAB. A fullword address in the PDAB (PDADCBEP) points to the address of the DCB. It should be noted that this pointer may be invalid from the time a CLOSE macro is issued to the issuing of the next parallel GET macro.

In Figure 26, not more than three data sets (MAXDCB=3 in the PDAB operand) will be open for parallel processing at any given time. Assuming that data definition statements and data sets are supplied, DATASET1, DATASET2, and DATASET3 will be opened for parallel input processing as specified in the input processing OPEN macro instruction. Other attributes of each data set are QSAM (MACRF=G), simple buffering by default, locate or move mode (MACRF=L or M), fixed length records (RECFM=F), and exit list entry for a PDAB (X'92'). Note that both locate and move modes may be used in the same data set queue. The mapping macros, DCBD and PDABD, are used to reference the DCBs and the PDAB respectively.

```
          . . .
          OPEN     (DATASET1,(INPUT),DATASET2,(INPUT),DATASET3,          X
                   (INPUT),DATASET4,(OUTPUT))
          TM       DATASET1+DCBQSWS-IHADCB,DCBPOPEN    Opened for parallel processing
          BZ       SEQRTN                              Branch on no to sequential routine
          TM       DATASET2+DCBQSWS-IHADCB,DCBPOPEN
          BZ       SEQRTN
          TM       DATASET3+DCBQSWS-IHADCB,DCBPOPEN
          BZ       SEQRTN
GETRTN    GET      DCBQUEUE,BUFFERAD,TYPE=P
          LR       10,1                      Save record pointer
          . . .
          . . .                              Record updated in place
          . . .
          PUT      DATASET4,(10)
          B        GETRTN
EODRTN    EQU      *                         Close DCB which just reached EODAD
          L        2,DCBQUEUE+PDADCBEP-IHAPDAB
          L        2,0(0,2)
          CLOSE    ((2))
          CLC      ZEROS(2),DCBQUEUE+PDANODCB-IHAPDAB      Any DCBs left?
          BL       GETRTN                                  Branch if yes
          . . .
          . . .
DATASET1  DCB      DDNAME=DDNAME1,DSORG=PS,MACRF=GL,RECFM=FB,           X
                   LRECL=80,EODAD=EODRTN,EXLST=SET3XLST
DATASET2  DCB      DDNAME=DDNAME2,DSORG=PS,MACRF=GL,RECFM=FB,           X
                   LRECL=80,EODAD=EODRTN,EXLST=SET3XLST
DATASET3  DCB      DDNAME=DDNAME3,DSORG=PS,MACRF=GMC,RECFM=FB,          X
                   LRECL=80,EODAD=EODRTN,EXLST=SET3XLST
DATASET4  DCB      DDNAME=DDNAME4,DSORG=PS,MACRF=PM,RECFM=FB,           X
                   LRECL=80
DCBQUEUE  PDAB     MAXDCB=3
SET3XLST  DC       0F'0', X'92', AL3(DCBQUEUE)
ZEROS     DC       X'0000'
          DCBD     DSORG=QS
          PDABD
          . . .
```

**Note:** The number of bytes required for PDAB is equal to 24+8n where n is the value of the keyword, MAXDCB.

Figure 26. Parallel Processing of Three Data Sets

Following the OPEN macro instruction, tests are made to determine whether the DCBs were opened for parallel processing. If not, the sequential processing routine is given control.

When one or more data sets are opened for parallel processing, the get routine retrieves a record, saves the pointer in register 10, processes the record, and writes it to DATASET4. This process continues until an end-of-data condition is detected on one of the input data sets; the end-of-data routine locates the completed input data set and removes it from the queue with the CLOSE macro instruction. A test is then made to determine whether any data sets remain on the queue. Processing continues in this manner until the queue is empty.

## Basic Access Technique

The basic access technique provides the READ and WRITE macro instructions for transmitting data between virtual and auxiliary storage. This technique is used when the operating system cannot predict the sequence in which the records are to be processed or when you do not want some or all of the automatic functions performed by the queued access technique. Although the system does not provide anticipatory buffering or synchronized scheduling, macro instructions are provided to help you program these operations.

The READ and WRITE macro instructions process blocks, not records. Thus, blocking and deblocking of records is your responsibility. Buffers, allocated by either you or the operating system, are filled or emptied individually each time a READ or WRITE macro instruction is issued. Moreover, the READ and WRITE macro instructions only initiate input/output operations. To ensure that the operation is completed successfully, you must issue a CHECK macro instruction to test the data event control block (DECB) or issue a WAIT macro instruction and then check the DECB yourself. (The only exception to this is when the SYNAD or EODAD routine is entered, neither a WAIT or CHECK macro instruction should be issued to previously outstanding READ or WRITE requests.) The number of READ or WRITE macro instructions issued before a CHECK macro instruction is used should not exceed the specified number of channel programs (NCP).

**Grouping Related Control Blocks in a Paging Environment:** In an OS/VS system, related control blocks (the DCB and DECB) and data areas (buffers and key areas) should be coded so they assemble in the same area of your program. This will reduce the number of paging operations required to read from and write to your data set.

**Using Overlapped I/O with BSAM:** When using BSAM with overlapped I/O (multiple I/O requests outstanding at one time), more than one DECB must be used. A different DECB should be specified for each channel program. For example, if you specify NCP=3 in your DCB for the data set and you are reading records from the data set, you should code the following macros in your program:

```
...
READ DECB1,...
READ DECB2,...
READ DECB3,...
CHECK DECB1
CHECK DECB2
CHECK DECB3
...
```

## READ—Read a Block

The READ macro instruction retrieves a data block from an input data set and places it in a designated area of virtual storage. To allow overlap of the input operation with processing, the system returns control to your program before the read operation is completed. The DECB created for the read operation must be tested for successful completion before the record is processed or the DECB is reused.

If an indexed sequential data set is being read, the block is brought into virtual storage and the address of the record is returned to you in the DECB.

When you use the READ macro instruction for BSAM to read a direct data set with spanned records and keys and you specify BFTEK=R in your DCB, the data management routines displace record segments after the first in a record by key length. Thus, you can expect the block descriptor word and the segment descriptor word at the same locations in your buffer or buffers, regardless of whether you read the first segment of a record, which is preceded in the buffer by its key, or a subsequent segment, which does not have a key. This procedure is called *offset reading*.

You can specify variations of the READ macro instruction according to the organization of the data set being processed and the type of processing to be done by the system as follows:

**Sequential**

| | | |
|---|---|---|
| SF | - | Read the data set sequentially. |
| SB | - | Read the data set backward (magnetic tape, format-F and format-U only). When RECFM=FBS, data sets with the last block truncated cannot be read backward. |

**Indexed Sequential**

| | | |
|---|---|---|
| K | - | Read the data set. |
| KU | - | Read for update. The system maintains the device address of the record; thus, when a WRITE macro instruction returns the record, no index search is required. |

**Direct**

| | | |
|---|---|---|
| D | - | Use the direct access method. |
| I | - | Locate the block using a block identification. |
| K | - | Locate the block using a key. |
| F | - | Provide device position feedback. |
| X | - | Maintain exclusive control of the block. |
| R | - | Provide next address feedback. |
| U | - | Next address can be a capacity record or logical record, whichever occurred first. |

## WRITE—Write a Block

The WRITE macro instruction places a data block in an output data set from a designated area of virtual storage. The WRITE macro instruction can also be used to return an updated record to a data set. To allow overlap of output operations with processing, the system returns control to your program before the write operation is completed. The DECB created for the write operation must be tested for successful completion before the DECB can be reused. For ASCII tape data sets, do not issue more than one WRITE on the same record, because the WRITE macro instruction causes the data in the record area to be translated from EBCDIC to ASCII.

As with the READ macro instruction, you can specify variations of the WRITE macro instruction according to the organization of the data set and the type of processing to be done by the system as follows:

**Sequential**

SF - Write the data set sequentially.

SFR - Write the data set sequentially with next-address feedback.

**Indexed Sequential**

K - Write a block containing an updated record, or replace a record with a fixed, unblocked record having the same key. The record to be replaced need not have been read into virtual storage.

KN - Write a new record or change the length of a variable-length record.

**Direct**

SD - Write a dummy fixed-length record.

SZ - Write a capacity record (R0). The system supplies the data, writes the capacity record, and advances to the next track.

D - Use the direct access method.

I - Search argument identifies a block.

K - Search argument is a key.

A - Add a new block.

F - Provide record location data (feedback).

X - Release exclusive control.

## CHECK—Test Completion of Read or Write Operation

When processing a data set, you can test for completion of a READ or WRITE request by issuing a CHECK macro instruction. The system tests for errors and exceptional conditions in the data event control block (DECB). Successive CHECK macro instructions issued for the same data set must be issued in the same order as the associated READ and WRITE macro instructions.

The check routine passes control to the appropriate exit routines specified in the DCB for error analysis (SYNAD) or, for sequential data sets, end-of-data (EODAD). It also automatically initiates end-of-volume procedures (volume switching or extending output data sets).

If you specify OPTCD=Q in the DCB, CHECK causes input data to be translated from ASCII to EBCDIC.

## WAIT—Wait for Completion of a Read or Write Operation

When processing a data set, you can test for completion of any READ or WRITE request by issuing a WAIT macro instruction. The input/output operation is synchronized with processing, but the DECB is not checked for errors or exceptional conditions, nor are end-of-volume procedures initiated. Your program must perform these operations.

For BDAM and BISAM, a WAIT macro must be issued for each READ or WRITE macro if MACRF=C is not coded in the associated DCB. When MACRF=C is coded, and at all times for BSAM and BPAM, a CHECK macro must be issued for each READ or WRITE macro. Since the CHECK macro incorporates the function of the WAIT macro, a WAIT is normally redundent for those access methods. The ECBLIST form of the WAIT macro may be useful, though, in selecting which of a number of outstanding events should be checked first.

The WAIT macro instruction can be used to await completion of multiple read and write operations. Each operation must then be checked or tested separately.

**Example:** You have opened an input DCB for BSAM with NCP=2, and an output DCB for BISAM with NCP=1 and without specifying MACRF=C. You have issued two BSAM READ macros and one BISAM WRITE macro. You now issue the WAIT macro with ECBLIST pointing to the BISAM DECB and the first BSAM DECB. (Since BSAM

requests are serialized, the first request must execute before the second one.) When you regain control, you will inspect the DECBs to see which has completed (second bit on). If it was BISAM, you will issue another WRITE macro. If it was BSAM, you will issue a CHECK macro and then another READ macro.

## Data Event Control Block (DECB)

A data event control block is a 16- to 32-byte area reserved by each READ or WRITE macro instruction. It contains control information and pointers to standard status indicators. It is described in detail in Appendix A of *OS/VS1 Data Management Macro Instructions*.

The DECB is examined by the check routine when the I/O operation is completed, to determine if an uncorrectable error or exceptional condition exists. If it does, control is passed to your SYNAD routine. If you have no SYNAD routine, the task is abnormally terminated.

## *Error Handling*

The basic and queued access techniques both provide special macro instructions for analyzing input/output errors. These macro instructions can be used in SYNAD routines and in error analysis routines that are entered directly when you use the basic access technique with indexed sequential data sets.

## SYNADAF—Perform SYNAD Analysis Function

The SYNADAF macro instruction analyzes the status, sense, and exceptional condition code data that is available to your error analysis routine. It produces an error message that your routine can write into any appropriate data set. The message is in the form of an unblocked variable-length record, but you can write it as a fixed-length record by omitting the block length and record length fields that precede the message text.

The text of the message is 120 characters long, and begins with a field of 36 or 42 blanks; you can use the blank field to add your own remarks to the message. Following is a typical message with the blank field omitted:

```
,TESTJOB฿,STEP2฿฿฿,283,TA,MASTER฿฿,READ฿,DATACHECK฿฿฿฿฿,
0000015,BSAM
```

This message indicates that a data check occurred during reading of the fifteenth block of a data set. The data set was identified by a DD statement named MASTER, and was on a magnetic-tape volume on unit 283. The name of the job was TESTJOB; the name of the job step was STEP2.

If the error analysis routine is entered because of an input error, the first 6 bytes of the message (bytes 8-13) contain binary information. If no data was transmitted or if the access method is QISAM, the first 6 bytes are blanks or binary zeros. If the error did not prevent data transmission, the first 6 bytes contain the address of the input buffer and the number of bytes read. You can use this information to process records from the block; for example, you might print each record after printing the error message. Before printing the message, however, you should replace this binary information with EBCDIC characters.

The SYNADAF macro instruction provides its own save area and makes this area available to your error analysis routine. When used at the entry point of a SYNAD routine, it fulfills the routine's responsibility for providing a save area.

**SYNADRLS—Release SYNADAF Message and Save Areas**

The SYNADRLS macro instruction releases the message and save areas provided by the SYNADAF macro instruction. You must issue this macro instruction before returning from the error analysis routine.

**ATLAS—Perform Alternate Track Location Assignment**

The ATLAS macro instruction enables your program to recover from permanent input/output errors when processing a data set in direct-access storage. After a data check, or in certain missing-address-marker conditions, you can issue ATLAS to assign an alternate track to replace the error track or to transfer data from the error track to the alternate track.

The use of this macro requires a knowledge of channel programming. A detailed description of the macro instruction and its use is included in *OS/VS1 Data Management for System Programmers*.

If you do not use the ATLAS macro instruction, you can use the IEHATLAS utility program to perform the same function. The principal difference between the macro instruction and the utility program is that the latter provides error recovery only after your own program has been completed. For a detailed description of IEHATLAS, refer to *OS/VS Utilities*.

## *Selecting an Access Method*

Access methods are identified primarily by the data set organization to which they apply. For instance, BDAM is the basic access method for direct organization. Nevertheless, there are times when an access method identified with one organization can be used to process a data set usually thought of as organized in a different manner. Thus, a data set created by the basic access method for sequential organization (BSAM) may be processed by the basic direct access method (BDAM). If the queued access technique is used to process a sequential data set, the access method is referred to as the queued sequential access method (QSAM).

Basic access methods are used for all data organizations, while queued access methods apply only to sequential and indexed sequential data sets as shown in Figure 27.

| Data Set Organization | Access Technique | |
|---|---|---|
| | Basic | Queued |
| Sequential | BSAM | QSAM |
| Partitioned | BPAM | |
| Indexed Sequential | BISAM | QISAM |
| Direct | BDAM | |

Figure 27. Data Management Access Methods

It is possible to control an I/O device directly while processing a data set with any data organization without using a specific access method. The execute channel program (EXCP) macro instruction uses the system programs that provide for scheduling and queuing I/O requests, efficient use of channels and devices, data protection, interruption procedures, error recognition and retry. Complete details about the EXCP macro are in *OS/VS1 Data Management for System Programmers*.

## *Opening and Closing a Data Set*

Although your program has been assembled, the various data management routines required for I/O operations are not a part of the object code. In other words, your program is not completely assembled until the DCBs are initialized for execution. You accomplish initialization by issuing the OPEN macro instruction. After all DCBs have been completed, the system ensures that all required access method routines are loaded and ready for use and that all channel command word lists and buffer areas are ready.

Access method routines are selected and loaded according to data control fields that indicate:

- Data organization

- Buffering technique

- Access technique

- I/O unit characteristics

This information is used by the system to allocate virtual-storage space and load the appropriate routines. These routines, the channel command word (CCW) lists, and buffer areas created automatically by the system remain in virtual storage until the close routine signals that they are no longer needed by the DCB that was using them.

When I/O operations for a data set are completed, you should issue a CLOSE macro instruction to return the DCB to its original status, handle volume disposition, create data set labels, complete writing of queued output buffers, and free virtual and auxiliary storage.

**Managing Buffer Pools when Closing Data Sets:** After closing the data set, you should issue a FREEPOOL macro instruction to release the virtual storage used for the buffer pool. If you plan to process other data sets, use FREEPOOL to regain the buffer pool storage space. If you expect to reopen a data set using the same DCB, use FREEPOOL unless the buffer pool created the first time the data set was opened will meet your needs when you reopen the data set. FREEPOOL is discussed in more detail in the section "Buffer Pool Construction."

After the data set has been closed, the DCB can be used for another data set. If you do not close the data set before a task terminates, the operating system closes it automatically. If the DCB is not available to the system at that time, the operating system abnormally terminates the task, and data results can be unpredictable. Note, however, that the operating system cannot automatically close any open data sets after the normal termination of a program that was brought into virtual storage by the loader. Therefore, loaded programs must include CLOSE macro instructions for all open data sets.

**Simultaneous Opening and Closing of Multiple Data Sets:** An OPEN or CLOSE macro instruction can be used to initiate or terminate processing of more than one data set. Simultaneous opening or closing is faster than issuing separate macro instructions; however, additional storage space is required for each data set specified. The coding examples in Figures 28 and 29 show the macro expansions for simultaneous open and close operations.

**Opening and Closing Data Sets Shared by More Than One Task:** When more than one task is sharing a data set, the following restrictions must be recognized. Failure to adhere to these restrictions endangers the integrity of the shared data set.

- All tasks sharing a DCB must be in the job step that opened the DCB (see "Sharing a Data Set").

- Each task sharing a DCB must ensure that all of the input and output operations it initiated using a given DCB are complete, before the task terminates. A CLOSE macro instruction issued for the DCB will ensure termination of all input and output operations.

- A DCB can be closed only by the task that opened it.

**Considerations for Opening and Closing Data Sets:**

- Two or more DCBs should never be concurrently open for output to the same data set on a direct-access device, except with the basic indexed sequential access method (BISAM). Otherwise the end-of-file record written by CLOSE for one DCB may overlay data associated with another DCB.

- If one DCB is concurrently open for input and one for output to the same data set on a direct-access device, the input DCB may be unable to read what the output DCB wrote if the output DCB extended the data set.

- If you want to use the same DD statement for two or more DCBs, you cannot specify parameters for fields in the first DCB and then be assured of obtaining the default parameters for the same fields in any subsequent DCB using the same DD statement. This is true for both input and output and is especially important when you are using more than one access method. Any action on one DCB that alters the JFCB affects the other DCB(s) and thus can cause unpredictable results. Therefore, unless the parameters of all DCBs using one DD statement are the same, you should use separate DD statements.

- Associated data sets for the 3525 Card Punch can be opened in any order, but all data sets must be opened before any processing can begin. Associated data sets can be closed in any order, but once a data set has been closed, I/O operations cannot be performed on any of the associated data sets. See *OS and OS/VS Programming Support for the IBM 3505 Card Reader and IBM 3525 Card Punch* for more information.

- Volume disposition specified in the OPEN or CLOSE macro instruction can be overridden by the system if necessary. However, you need not be concerned; the system automatically requests the mounting and demounting of volumes, depending upon the availability of devices at a particular time. Additional information on volume disposition is provided in *OS/VS1 JCL Reference.*

- Two or more DCBs should not be concurrently open using the same ddname. Any action on one DCB that alters the task input/output table (TIOT) or JFCB affects the other DCB(s) and thus can cause unpredictable results. However, a concurrent open of a single volume data set allocated to a single unit does not adversely affect the TIOT. Thus, a one volume physical sequential DASD data set can be multiply open for input, or within the other restrictions, for output. A multivolume ISAM or BDAM data set can be concurrently open as long as the volume serial list remains the same.

  If an abnormal termination occurs during open, close, or end-of-volume processing, the DCB causing the error and any DCBs being processed in parallel with it may be left in such a state that they cannot be used and cannot be closed (or will be closed incorrectly). The device associated with the DCB in error (such as a magnetic tape

unit) can become locked to this DCB and will not be reuseable until the end of the job step when device deallocation will unlock it. This is true whether or not the abnormal termination is intercepted by a STAE routine, a DCB ABEND exit routine, or a mother task. When opening or closing DCBs in parallel, specifying delay or ignore abnormal termination in the DCB ABEND exit routine will allow the other DCBs to complete processing normally.

- During task termination, the system issues a CLOSE macro for each data set which is still open. If this is an abnormal termination, the QSAM close routines (which would normally finish processing buffers) are bypassed. Any outstanding I/O requests are purged. Thus, your last data records may be lost for a QSAM output data set.

- It is a good procedure to close an ISAM data set before task termination because, if an I/O error is detected, the ISAM close routines cannot return the problem program registers to the SYNAD routine, causing unpredictable results.

## OPEN—Prepare a Data Set for Processing

The OPEN macro instruction is used to complete a data control block for an associated data set. The method of processing and the volume positioning instruction in the event of an end-of-volume condition can be specified.

**Processing Method:** You can process a data set as either input or output (by coding INPUT, OUTPUT, or EXTEND as the processing method operand of the OPEN macro) or, under BSAM, by coding INOUT, OUTIN, or OUTINX. If the data set resides on a direct-access volume, you can code UPDAT in the processing method operand to indicate that records can be updated. By coding RDBACK in this operand, you can specify that a magnetic-tape volume containing format-F or format-U records is to be read backward. Variable-length records cannot be read backward. If the processing method operand is omitted from the OPEN macro instruction, INPUT is assumed. The operand is ignored by the basic indexed sequential access method (BISAM); it must be specified as OUTPUT or EXTEND when you are using the queued indexed sequential access method (QISAM) to create an indexed sequential data set. You can override the INOUT, OUTIN, or OUTINX at execution by using the LABEL parameter of the DD statement, as discussed in *OS/VS1 JCL Reference*.

SYSIN and SYSOUT data sets must be opened for INPUT and OUTPUT, respectively. INOUT is treated as INPUT; OUTIN, EXTEND, or OUTINX are treated as OUTPUT. UPDAT and RDBACK cannot be used.

In Figure 28, the data sets associated with three DCBs are to be opened simultaneously.

```
        OPEN    ( TEXTDCB, ,CONVDCB,( OUTPUT ),PRINTDCB,( OUTPUT ) )
+       CNOP    0,4                 Align list to fullword
+       BAL     1,*+16              Load reg1 w/list address
+       DC      AL1( 0 )            Option byte
+       DC      AL3( TEXTDCB )      DCB address
+       DC      AL1( 15 )           Option byte
+       DC      AL3( CONVDCB )      DCB address
+       DC      AL1( 143 )          Option byte
+       DC      AL3( PRINTDCB )     DCB address
+       SVC     19                  Issue open SVC
```
Figure 28. Opening Three Data Sets Simultaneously

Since no processing method operand is specified for TEXTDCB, the system assumes INPUT. Both CONVDCB and PRINTDCB are opened for output. No volume positioning options are specified; thus, the disposition indicated by the DD statement DISP parameter is used.

At execution, the SVC 19 instruction passes control to the Open routine, which then initializes the three DCBs and loads the appropriate access method routines.

## CLOSE—Terminate Processing of a Data Set

The CLOSE macro instruction is used to terminate processing of a data set and release it from a DCB. The volume positioning that is to result from closing the data set can also be specified. Volume positioning options are the same as those that can be specified for end-of-volume conditions in the OPEN macro instruction or the DD statement. An additional volume positioning option, REWIND, is available and can be specified by the CLOSE macro instruction for magnetic-tape volumes. REWIND positions the tape at the load point regardless of the direction of processing.

You can code CLOSE TYPE=T and perform some close functions for sequential data sets on magnetic tape and direct-access volumes processed with BSAM. When you use TYPE=T, the DCB used to process the data set maintains its open status, and you should not issue another OPEN macro instruction to continue processing the same data set. This option cannot be used in a SYNAD routine.

The TYPE=T operand causes the system control program to process labels, modify some of the fields in the system control blocks for that data set, and reposition the volume (or *current* volume in the case of multivolume data sets) in much the same way that the normal CLOSE macro does. When you code TYPE=T, you can specify that the volume either be positioned at the end of data (the LEAVE option) or be repositioned at the beginning of data (the REREAD option). Magnetic-tape volumes are repositioned either immediately before the first data record or immediately after the last data record; the presence of tape labels has no effect on repositioning. Figure 29, which assumes a sample data set containing 1000 records, illustrates the relationship between each positioning option and the point at which you resume processing the data set after issuing the temporary close.

Begin processing data set →

Begin processing tape data set (open for read backward) ↘

| Record 1 | Record 2 | Record 3 | | Record 999 | Record 1000 |

| If you CLOSE TYPE = T and specify | After temporary close, you will resume processing |
|---|---|
| LEAVE | Immediately after record 1000 |
| LEAVE (with tape data set open for read backward) | Immediately before record 1 |
| REREAD | Immediately before record 1 |
| REREAD (with tape data set open for read backward) | Immediately after record 1000 |

Figure 29. Record Processed When LEAVE or REREAD is Specified for CLOSE TYPE=T

If you code the release (RLSE) operand on the DD statement for the data set, it is ignored by temporary close, but any unused space will be released when you finally issue the normal CLOSE macro instruction.

It is possible to use BSAM to process a data set that is not physical-sequential; if you use CLOSE TYPE=T for them, the following restrictions apply:

- The DCB for the data set you are processing on a direct-access device must specify either DSORG=PS or DSORG=PSU for input processing, and either DSORG=PS, DSORG=PSU, DSORG=PO, or DSORG=POU for output processing.

- The DCB must not be open for input to a member of a partitioned data set.

- If you open a data set on a direct-access device for output and issue CLOSE TYPE=T, the volume will be repositioned only if the data set was created with DSORG=PS, DSORG=PSU, DSORG=PO, or DSORG=POU (you cannot specify the REREAD option if DSORG=PO or DSORG=POU is specified). (This restriction prohibits the use of temporary close following or during the building of a BDAM data set that is created by specifying BSAM MACRF=WL).

- If you open the data set for input and issue CLOSE TYPE=T with the LEAVE option, the volume will be repositioned only if the data set specifies DSORG=PS or DSORG=PO.

**Note:** When a data control block is shared among multiple tasks, only the task that opened the data set can close it unless TYPE=T is specified.

Before issuing the CLOSE macro, a CHECK macro must be issued for all DECBs that have outstanding I/O from WRITE macro instructions. When CLOSE TYPE=T is specified, a CHECK macro must be issued for all DECBs that have outstanding I/O from either WRITE or READ macro instructions.

In Figure 30, the data sets associated with three DCBs are to be closed simultaneously.

| | CLOSE | ( TEXTDCB , , CONVDCB , , PRINTDCB ) | |
|---|---|---|---|
| + | CNOP | 0,4 | Align list to fullword |
| + | BAL | 1,*+16 | Load reg1 w/list addr |
| + | DC | AL1(0) | Option byte |
| + | DC | AL3(TEXTDCB) | DCB address |
| + | DC | AL1(0) | Option byte |
| + | DC | AL3(CONVDCB) | DCB address |
| + | DC | AL1(128) | Option byte |
| + | DC | AL3(PRINTDCB) | DCB address |
| + | SVC | 20, | Issue close SVC |

Figure 30. Closing Three Data Sets Simultaneously

Because no volume positioning operands are specified, the position indicated by the DD statement DISP parameter is used.

At execution, the SVC 20 instruction passes control to the Close routine, which terminates processing of the three data sets and returns the three DCBs to their original status.

Control is passed automatically to the data management end-of-volume routine when any of the following conditions is detected:

- Tapemark (input tape volume)

- Filemark or end of last extent (input direct-access volume)

- End-of-data indicator (input device other than magnetic tape or direct-access volume). An example of this would be the last card read on a card reader.

- End of reel (output tape volume)

- End of extent (output direct-access volume)

You may issue a force end-of-volume (FEOV) macro instruction before the end-of-volume condition is detected.

The end-of-volume routine checks or creates standard trailer labels, if the LABEL parameter of the associated DD statement indicates standard labels. If UL or NSL is specified, control is passed to the appropriate user label routine (if it is specified in your exit list.)

If multiple-volume data sets are specified in your DD statement, automatic volume switching is accomplished by the end-of-volume routine. When an end-of-volume condition exists on an output data set, additional space is allocated as indicated in your DD statement. If no more volumes are specified or if more than specified are required, the storage is obtained from any available volume on a device of the same type. If no such volume is available, your job is terminated.

**Volume Positioning:** When an end-of-volume condition is detected, the system positions the volume according to the disposition specified in the DD statement unless the volume disposition is specified in the OPEN macro instruction. Volume positioning instructions for a sequential data set on magnetic tape can be specified as LEAVE or REREAD.

LEAVE
positions a labeled tape to the point following the tape mark that follows the data set trailer label group, and an unlabeled volume to the point following the tape mark that follows the last block of the data set.

REREAD
positions a labeled tape to the point preceding the data set header label group, and an unlabeled tape to the point preceding the first block of the data set. If the tape was last read backward, LEAVE and REREAD do the following:

LEAVE
positions a labeled tape to the point preceding the data set header label group, and an unlabeled tape to the point preceding the first block of the data set.

REREAD
positions a labeled tape to the point following the tape mark that follows the data set trailer label group, and an unlabeled tape to the point following the tape mark that follows the last block of the data set.

If, however, you want to position the current volume according to the option specified in the DISP parameter of the DD statement, you code DISP in the OPEN macro instruction.

DISP

specifies that a tape volume is to be disposed of in the manner implied by the DD statement associated with the data set. Direct-access volume positioning and disposition are not affected by this parameter of the OPEN macro instruction. There are several dispositions that can be specified in the DISP parameter of the DD statement; DISP can be PASS, DELETE, KEEP, CATLG, or UNCATLG.

If the number of volumes is less than or equal to the number of allocated units, the current volume will be positioned as follows:

| DISP Parameter | Action |
| --- | --- |
| PASS | Forward space to the end of data set on the current volume. |
| DELETE | Rewind the current volume. |
| KEEP, CATLG, or UNCATLG | Rewind and unload the current volume. |

If the number of volumes is greater than the number of units allocated, the current volume will be rewound and unloaded.

A volume positioning instruction can be specified only if the processing method operand has been specified. It is ignored if devices other than magnetic-tape and direct-access are used, or if the number of volumes exceeds the number of available units.

For magnetic-tape volumes that are not being unloaded, positioning varies according to the direction of the last input operation and the existence of tape labels. If the tape was last read forward:

LEAVE

positions a labeled tape to the point following the tapemark that follows the data set trailer label group, and an unlabeled volume to the point following the tapemark that follows the last block of the data set.

REREAD

positions a labeled tape to the point preceding the data set header label group, and an unlabeled tape to the point preceding the first block of the data set.

If the tape was last read backward:

LEAVE

positions a labeled tape to the point preceding the data set header label group, and an unlabeled tape to the point preceding the first block of the data set.

REREAD

positions a labeled tape to the point following the tapemark that follows the data set trailer label group, and an unlabeled tape to the point following the tapemark that follows the last block of the data set.

## FEOV—Force End of Volume

The FEOV macro instruction directs the operating system to initiate end-of-volume processing before the physical end of the current volume is reached. If another volume has been specified for the data set, volume switching takes place automatically. The volume positioning options REWIND and LEAVE are available.

If an FEOV macro is issued for a spanned multivolume data set which is being read using QSAM, errors may occur when the next GET macro is issued. These errors are documented in the section, "Spanned Variable-Length Records" in "Part 1: Introduction to Data Management."

The FEOV macro instruction can only be used when you are using BSAM or QSAM. FEOV is ignored if issued for a SYSIN or SYSOUT data set.

# Buffer Acquisition and Control

The operating system provides several methods of buffer acquisition and control. Each *buffer* (virtual-storage area used for intermediate storage of input/output data) usually corresponds in length to the size of a block in the data set being processed. When you use the queued access technique, any reference to a buffer actually refers to the next record (*buffer segment*).

You can assign more than one buffer to a data set by associating the buffer with a *buffer pool*. A buffer pool must be constructed in a virtual-storage area allocated for a given number of buffers of a given length.

The number of buffers you assign to a data set should be a tradeoff against the frequency with which you refer to each buffer. A buffer that is not referred to for a relatively long period of time may be paged out. If this were allowed to happen to any considerable degree, it could result in a greater number of buffers actually decreasing throughout.

Buffer segments and buffers within the buffer pool are controlled automatically by the system when the queued access technique is used. However, you can terminate processing of a buffer by issuing a release (RELSE) macro instruction for input or a truncate (TRUNC) macro instruction for output. Two buffering techniques, simple and exchange, can be used to process a sequential data set. Only simple buffering can be used to process an indexed sequential data set.

If you use the basic access technique, you can use buffers as work areas rather than as intermediate storage areas. You can control them directly, by using the GETBUF and FREEBUF macro instructions, or dynamically for BDAM and BISAM, by requesting dynamic buffering in your DCB macro instruction and your READ or WRITE macro instruction. If you request dynamic buffering, the system will automatically provide a buffer each time a READ macro instruction is issued. That buffer will be freed when you issue a WRITE or FREEDBUF macro instruction.

## Buffer Pool Construction

Buffer pool construction can be accomplished in any of three ways:

- Statically using the BUILD macro instruction
- Explicitly using the GETPOOL macro instruction
- Automatically by the system when the data set is opened

If QSAM simple buffering is used, the buffers are automatically returned to the pool when the data set is closed. If the buffer pool is constructed explicitly or automatically, the virtual storage area must be returned to the system by the FREEPOOL macro instruction.

In many applications, fullword or doubleword alignment of a block within a buffer is important. You can specify in the DCB that buffers are to start on either a doubleword boundary or a fullword boundary that is not also a doubleword boundary (by coding BFALN=D or F). If doubleword alignment is specified for format-V records, the fifth byte of the first record in the block is so aligned. For that reason, fullword alignment must be requested to align the first byte of the variable-length record on a doubleword boundary. The alignment of the records following the first in the block depends on the length of the previous records.

Note that buffer alignment provides alignment only for the buffer. If records from ASCII magnetic tape are read and the records use the block prefix, the boundary alignment of logical records within the buffer depends on the length of the block prefix. If the length is 4, logical records are on fullword boundaries. If the length is 8, logical records are on doubleword boundaries.

If the BUILD macro instruction is used to construct the buffer pool, alignment depends on the alignment of the first byte of the reserved storage area.

When you process multiple QISAM data sets, you can use a common buffer pool. To do this, however, you must use the BUILD macro instruction to reformat the buffer pool before opening each data set.

## BUILD—Construct a Buffer Pool

When you know, before program assembly, both the number and the size of the buffers required for a given data set, you can reserve an area of appropriate size to be used as a buffer pool. Any type of area can be used—for example, a predefined storage area or an area of coding no longer needed.

A BUILD macro instruction, issued during execution of your program, structures the reserved storage area into a buffer pool. The address of the buffer pool must be the same as that specified for the buffer pool control block (BUFCB) in your DCB. The buffer pool control block is an 8-byte field preceding the buffers in the buffer pool. The number (BUFNO) and length (BUFL) of the buffers must also be specified. For QSAM, the length of BUFL must be at least the blocksize.

When the data set using the buffer pool is closed, you can reuse the area as required. You can also reissue the BUILD macro instruction to reconstruct the area into a new buffer pool to be used by another data set.

You can assign the buffer pool to two or more data sets that require buffers of the same length. To do this, you must construct an area large enough to accommodate the total number of buffers required at any one time during execution. That is, if each of two data sets requires five buffers (BUFNO=5), the BUILD macro instruction should specify ten buffers. The area must also be large enough to contain the 8-byte buffer pool control block.

## BUILDRCD—Build a Buffer Pool and a Record Area

The BUILDRCD macro instruction, like the BUILD macro instruction, causes a buffer pool to be constructed in an area of virtual storage you provide. In addition, BUILDRCD makes it possible for you to access variable-length, spanned records as complete logical records, rather than as segments.

You must be processing with QSAM in the locate mode and you must be processing either VS or VBS records, if you want to access the variable-length, spanned records as logical records. If you issue the BUILDRCD macro before the data set is opened, or during your DCB exit routine, you automatically get logical records rather than segments of spanned records.

Only one logical record storage area is built, no matter how many buffers are specified; therefore, you can't share the buffer pool with other data sets that may be open at the same time.

## GETPOOL—Get a Buffer Pool

If a specified area is not reserved for use as a buffer pool, or you want to defer specifying the number and length of the buffers until execution of your program, you should use the GETPOOL macro instruction. It enables you to vary the size and number of buffers according to the needs of the data set being processed.

The GETPOOL macro instruction structures a virtual-storage area allocated by the system into a buffer pool, assigns a buffer pool control block, and associates the pool with a specific data set. The GETPOOL macro instruction should be issued either before opening of the data set or during your DCB exit routine.

When using GETPOOL with QSAM, specify a buffer length (BUFL) of at least as large as the blocksize.

## Automatic Buffer Pool Construction

If you have requested a buffer pool and have not used an appropriate macro instruction by the end of your DCB exit routine, the system automatically allocates virtual-storage space for a buffer pool. The buffer pool control block is also assigned and the pool is associated with a specific DCB. For BSAM, a buffer pool is requested by specifying BUFNO. For QSAM, BUFNO can be specified or allowed to default to 2 (for all devices except unit-record) or 3 (for 2540) by the open routines. If you are using the basic access technique to process an indexed sequential or direct data set, you must indicate dynamic buffer control. Otherwise, the system does not construct the buffer pool automatically.

Because a buffer pool obtained automatically is not freed automatically when you issue a CLOSE macro instruction, you should also issue a FREEPOOL or FREEMAIN macro instruction, which is discussed in the next section.

## FREEPOOL—Free a Buffer Pool

Any buffer pool assigned to a DCB either automatically by the OPEN macro instruction (except when dynamic buffer control is used) or explicitly by the GETPOOL macro instruction should be released before your program is terminated. The FREEPOOL macro instruction should be issued to release the virtual-storage area as soon as the buffers are no longer needed. When you are using the queued access technique, a data set must be closed first. When you are using exchange buffering, the buffer pool must not be released until all the data sets have been closed.

If the OPEN macro was issued while running under a protect key of zero, a buffer pool which was obtained by OPEN should be released by issuing the FREEMAIN macro instead of the FREEPOOL macro. This is necessary because the buffer pool acquired under these conditions will be in storage assigned to subpool 252.

**Constructing a Buffer Pool:** Figures 31 and 32 illustrate several possible methods of constructing a buffer pool. They do not take into account the method of processing or controlling the buffers in the pool.

In Figure 31, a static storage area named INPOOL is allocated during program assembly. The BUILD macro instruction, issued during execution, arranges the buffer pool into ten buffers, each 52 bytes long. Five buffers are assigned to INDCB and five to OUTDCB, as specified in the DCB macro instruction for each. The two data sets share the buffer pool because both specify INPOOL as the buffer pool control block. Notice that an additional 8 bytes have been allocated for the buffer pool to contain the buffer pool control block.

```
          ...                              Processing
          BUILD     INPOOL,10,52                 Structure a buffer pool
          OPEN      (INDCB,,OUTDCB,(OUTPUT))
          ...                              Processing
ENDJOB    CLOSE     (INDCB,,OUTDCB)
          ...                              Processing
          RETURN                           Return to system control
INDCB     DCB       BUFNO=5,BUFCB=INPOOL,EODAD=ENDJOB,---
OUTDCB    DCB       BUFNO=5,BUFCB=INPOOL,---
          CNOP      0,8                          Force boundary alignment
INPOOL    DS        CL528                        Buffer pool
          ...
```

Figure 31. Constructing a Buffer Pool From a Static Storage Area

In Figure 32, two buffer pools are constructed explicitly by the GETPOOL macro instructions. Ten input buffers are provided, each 52 bytes long, to contain one fixed-length record; five output buffers are provided, each 112 bytes long, to contain two blocked records plus an 8-byte count field (required by ISAM). Notice that both data sets are closed before the buffer pools are released by the FREEPOOL macro instructions. The same procedure should be used if the buffer pools were constructed automatically by the OPEN macro instruction.

## Buffer Control

Your program can use four techniques to control the buffers used by your program. The advantages of each depend to a great extent upon the type of job you are doing. Simple and exchange buffering are provided for the queued access technique. The basic access technique provides for either direct or dynamic buffer control.

Although only simple buffering can be used to process an indexed sequential data set, buffer segments and buffers within a buffer pool are controlled automatically by the operating system.

In addition, the queued access technique provides four processing modes that determine the extent of data movement in virtual storage. Move, data, locate, or substitute mode processing can be specified for either the GET or PUT macro instruction. The buffer processing mode is specified in the MACRF field of the DCB macro instruction. The movement of a record is determined as follows:

- *Move mode:* The record is moved from an input buffer to your work area, or from your work area to an output buffer.

- *Data mode (QSAM format-V spanned records only):* The same as the move mode except only the data portion of the record is moved.

- *Locate mode:* The record is not moved. Instead, the address of the next input or output buffer is placed in register 1. For QSAM format-V spanned records, if you have specified logical records by specifying BFTEK=A or by issuing the BUILDRCD macro instruction, the address returned in register 1 points to a record area where the spanned record is assembled or segmented.

The PUT-locate routine uses the value in the DCBLRECL field to determine whether another record will fit into your buffer. Therefore, when you write a short record, you can maximize the number of records per block by modifying the DCBLRECL field

```
          . . .
          GETPOOL    INDCB,10,52                      Construct a 10-buffer pool
          GETPOOL    OUTDCB,5,112                     Construct a 5-buffer pool
          OPEN       (INDCB,,OUTDCB,(OUTPUT))
          . . .
ENDJOB    CLOSE      (INDCB,,OUTDCB)
          FREEPOOL   INDCB                            Release buffer pools after all
                                                      I/O is complete
          FREEPOOL   OUTDCB
          . . .
          RETURN                                      Return to system control
INDCB     DCB        DSORG=PS,BFALN=F,LRECL=52,RECFM=F,EODAD=ENDJOB,---
OUTDCB    DCB        DSORG=IS,BFALN=D,LRECL=52,KEYLEN=10,BLKSIZE=104,        C
          . . .      RKP=0,RECFM=FB,---
```
Figure 32. Constructing a Buffer Pool Using GETPOOL and FREEPOOL

*before* you issue a PUT-locate to get a buffer segment for the short record. The processing sequence follows:

1. Register 1 is returned to you with the address of the next buffer segment.

2. Move the record into the output buffer segment.

3. Put the length of the *next* (short) record into DCBLRECL.

4. Issue PUT-locate.

5. Move the short record into the buffer segment.

- *Substitute mode:* The record is not moved. Instead, the address of the next input or output buffer is interchanged with the address of your work area.

Two processing modes of the PUTX macro instruction can be used in conjunction with a GET-locate macro instruction. The update mode returns an updated record to the data set from which it was read; the output mode transfers an updated record to an output data set. There is no actual movement of data in virtual storage. The processing mode is specified by the operand of the PUTX macro instruction, as explained in *OS/VS1 Data Management Macro Instructions.*

If you use the basic access technique, you can control buffers in one of two ways:

- Directly, using the GETBUF macro instruction to retrieve a buffer constructed as described above. A buffer can then be returned to the pool by the FREEBUF macro instruction.

- Dynamically, by requesting a dynamic buffer in your READ or WRITE macro instruction. This technique can be used only when you are using BISAM or BDAM. If you request dynamic buffering, the system automatically provides a buffer each time a READ macro instruction is issued. The buffer is supplied from a buffer pool that is created by the system when the data set is opened. The buffer is released (returned to the pool) upon completion of a WRITE macro instruction when you are updating. If you do not update the record in the buffer and thus release the buffer when the record is written, the FREEDBUF macro instruction may be used. If you are processing an indexed sequential data set, the buffer is automatically released upon the next issuance of the READ macro instruction if there has been no intervening WRITE or FREEDBUF macro instruction.

## Simple Buffering

The term *simple buffering* refers to the relationship of segments within the buffer. All segments in a simple buffer are together in storage and are always associated with the same data set. When the buffer pool is constructed, the system creates a channel command word (CCW) for each buffer in the buffer pool. For this reason, each record must be physically moved from an input buffer segment to an output buffer segment. It can be processed within either segment or in a work area.

If you use simple buffering, records of any format can be processed. New records can be inserted and old records deleted as required to create a new data set. A record can be moved and processed as follows:

- Processed in an input buffer and then moved to an output buffer (GET-locate, PUT-move/PUTX-output)

- Moved from an input buffer to an output buffer where it can be processed (GET-move, PUT-locate)

- Moved from an input buffer to a work area where it can be processed and then moved to an output buffer (GET-move, PUT-move)

- Processed in an input buffer and returned to the data set (GET-locate, PUTX-update)

The following examples illustrate the control of simple buffers and the processing modes that can be used. The buffer pools may have been constructed in any way previously described.

**Simple Buffering—GET-locate, PUT-move/PUTX-output:** The GET macro instruction (step A, Figure 33) locates the next input record to be processed. Its address is returned in register 1 by the system. The address is passed to the PUT macro instruction in register 0.

The PUT macro instruction (step B, Figure 33) specifies the address of the record in register 0. The system then moves the record to the next output buffer.

**Note:** The PUTX-output macro instruction can be used in place of the PUT-move macro instruction. However, processing will be as described under Exchange Buffering (see PUT-substitute).

**Simple Buffering—GET-move, PUT-locate:** The PUT macro instruction locates the address of the next available output buffer. Its address is returned in register 1 and is passed to the GET macro instruction in register 0.

The GET macro instruction specifies the address of the output buffer into which the system moves the next input record.

A filled output buffer is not written until the next PUT macro instruction is issued.

**Simple Buffering—GET-move, PUT-move:** The GET macro instruction (step A, Figure 34) specifies the address of a work area into which the system moves the next record from the input buffer.

The PUT macro instruction (step B, Figure 34) specifies the address of a work area from which the system moves the record into the next output buffer.

**Simple Buffering—GET-locate, PUT-locate:** The GET macro instruction (step A, Figure 35) locates the address of the next available input buffer. The address is returned in register 1.

The PUT macro instruction (step B, Figure 35) locates the address of the next available output buffer. Its address is returned in register 1. You must then move the record from
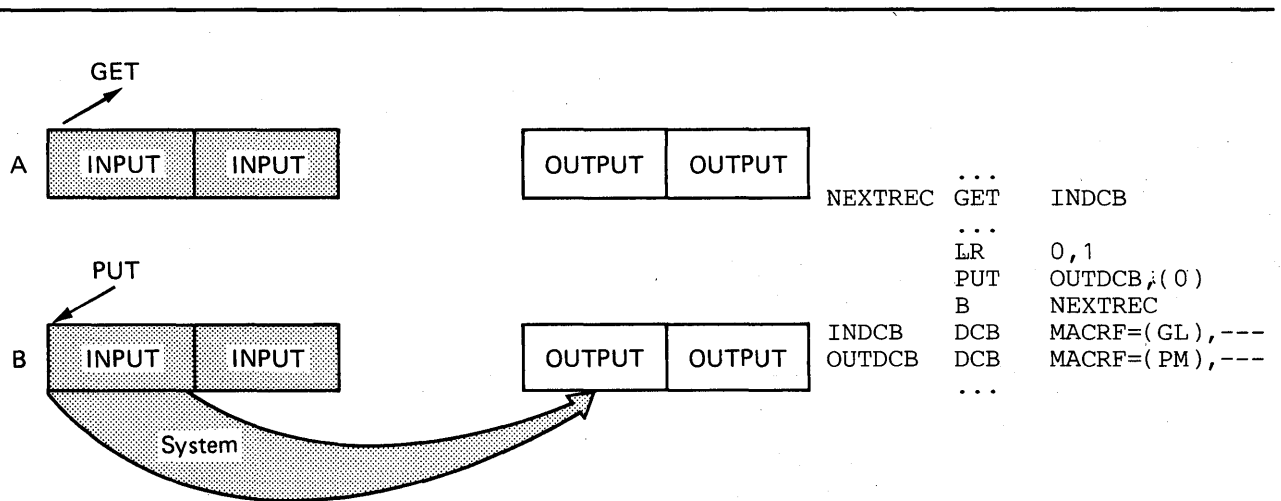


Figure 33. Simple Buffering with MACRF=GL and MACRF=PM

```
NEXTREC   GET   INDCB,WORKAREA
          ...
          PUT   OUTDCB,WORKAREA
          B     NEXTREC
WORKAREA  DS    CL50
INDCB     DCB   MACRF=(GM),---
OUTDCB    DCB   MACRF=(PM),---
          ...
```
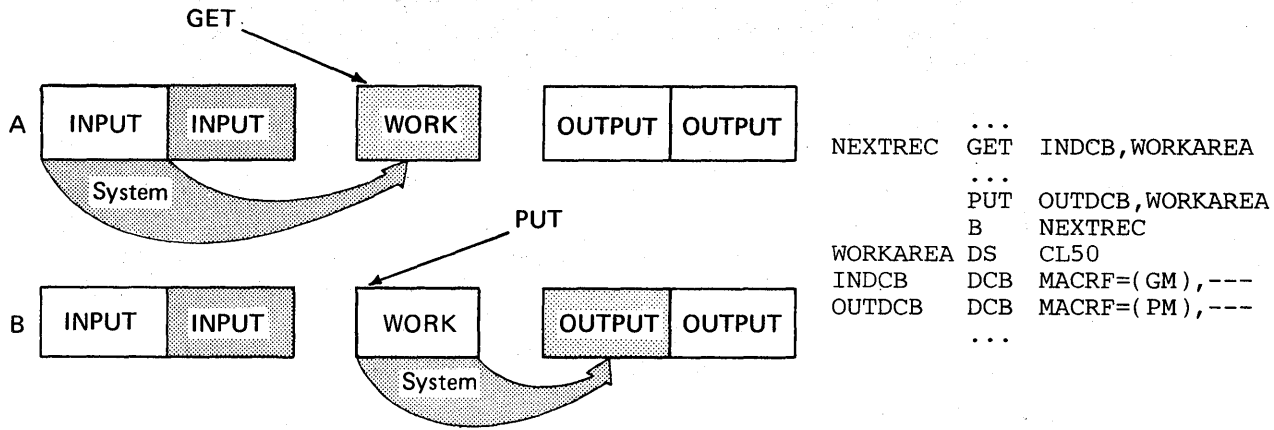
Figure 34. Simple Buffering with MACRF=GM and MACRF=PM



```
          ...
NEXTREC   GET    INDCB
          LR     7,1
          PUT    OUTDCB
          LR     6,1
          LA     5,INDCB
          USING  IHADCB,5
          LH     4,DCBLRECL
          SH     4,=H'1'
          EX     4,MOVEREC
          ...
          B      NEXTREC
MOVEREC   MVC    0(1,6),0(7)
INDCB     DCB    MACRF=(GL),
                 EODAD=EOF,---
OUTDCB    DCB    MACRF=(PL),---
          DCBD   DSORG=(LR)
EOF       ...
```

Figure 35. Simple Buffering with MACRF=GL and MACRF=PL

the input buffer to the output buffer (step C, Figure 35). Processing can be done either before or after the move operation.

A filled output buffer is not written until the next PUT macro instruction is issued. The CLOSE and FEOV macro instructions write the last record of your data set by issuing TRUNC and PUT macro instructions. Be careful not to issue an extra PUT before issuing CLOSE or FEOV. Otherwise, when the CLOSE or FEOV macro instruction tries to write your last record, the extra PUT will write a meaningless record or produce a sequence error.

Note that if records other than format-F records are being moved, the length attribute of the MVC instruction must be changed as shown by the code beginning with the USING statement in Figure 35. If the record is more than 256 bytes, you can code a move routine or use a MVCL instruction to process the complete record.

**Simple Buffering—UPDAT Mode:** When a data set is opened with UPDAT specified (Figure 36), only GET-locate and PUTX-update are supported. The GET macro locates the next input record to be processed and its address is returned in register 1 by the system. The user may update the record and issue a PUTX macro which will cause the block to be written back in its original location in the data set after all the logical records in that block have been processed.

## Exchange Buffering

The term *exchange buffering* refers to the relationship of segments within a buffer. All the segments in an exchange buffer are not necessarily together in virtual storage, nor are they always associated with the same data set. When the buffer pool is constructed, the system creates a channel command word (CCW) for each buffer segment in the buffer. The system can then combine all the segments (the logical records) of a buffer (a physical record or block of data) by using a technique called *data chaining*. In data chaining, a series of CCWs representing the buffer segments are chained and executed, thereby improving performance. However, when using exchange buffering, you must be sure that the CPU and channel can exchange CCWs faster than the I/O device can transfer data. Otherwise, as in the case of some of the faster direct-access storage devices used in combination with some of the slower CPUs, a hardware error called *data overrun* will result.

If a request for exchange buffering is denied by the system for any reason, move mode will be used instead. Move mode is compatible with exchange buffering.

To use exchange buffering, you must provide a work area comparable in size and alignment to a buffer segment. That work area is substituted for the next buffer segment (the storage areas change roles). The CCW created for the buffer segment actually points to the work area.

Why use exchange buffering? Because there is no need to move the record. This can mean a saving in processing time when you use substitute mode or PUTX-output mode. The use of exchange buffering during execution of your program requires these conditions:

- Input and output buffers must be of the same size and alignment.
- Records must be format-F records (blocked or unblocked).
- ASCII records must be format-F records with BUFOFF=(0).
- Track overflow cannot be used with blocked format-F records.
- GET-move and PUT-locate modes cannot be used.
- Unit-record devices must not be specified.

```
      ┌──► GET
      │ ┌─────────►
┌─────┴─┴──────┬───────────────┐
│   INPUT/     │    INPUT/      │
│   OUTPUT     │    OUTPUT·     │
└─────▲────────┴───────────────┘
      │  ┌──────────
      └──┴──► PUTX

(No movement of data takes place)
```

|  | . . . | |
|---|---|---|
|  | OPEN | ( UPDCB, ( UPDAT ) ) |
|  | . . . . | |
| NEXTREC | GET | UPDCB |
|  | . . . . | |
|  | PUTX | UPDCB |
|  | B | NEXTREC |
| UPDCB | DCB | MACRF=( GL,PM ),--- |
|  | . . . | |
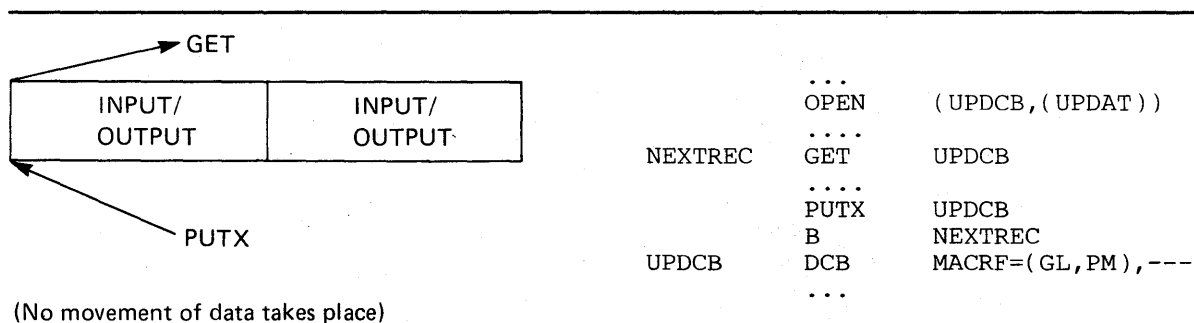
Figure 36. Simple Buffering with MACRF=GL and MACRF=PM-UPDAT Mode

If you request exchange buffering, but it cannot be used, the system automatically uses simple buffering. Move mode processing is used in place of substitute mode.

After opening the data set, you can test the DCBCIND1 field of the DCB to determine if simple buffering was substituted for exchange buffering because of inconsistencies in the DCB information. The eighth bit of the DCBCIND1 field is 1 for exchange buffering and 0 for simple buffering.

If your records are blocked format-F records, each segment is aligned as specified in the DCBBFALN field. Therefore, your buffer length (DCBBUFL) must specify buffers large enough to contain segments whose length is a multiple of 16 bytes. Otherwise, the specified boundary alignment cannot be achieved; simple buffering is used and only the first byte in the first record is aligned as specified.

To reopen a DCB that has been opened for exchange buffering, you must first close all DCBs using the buffer pool associated with the DCB to be reopened and issue a FREEPOOL macro instruction specifying the DCB to be reopened. There are two possible conditions that cannot be prechecked by the system:

- Word alignment that does not correspond to the characteristics of the machine. For maximum efficiency on a System/370, your record length should be a multiple of 8.

- An I/O device that transfers the data faster than the CPU can exchange the addresses in the CCW, which may cause data overrun. Several factors should be considered: (1) the relative rate of data transfer between the CPU and the I/O device, (2) the logical record length, and (3) the blocking factor. This problem can be avoided by making the logical record size equal to the block size to prevent data chaining of record segments. Since this problem is closely related to installation planning and system configuration, you may want to consult your installation planning representative if it is a significant problem.

**Note:** Performance can be adversely affected when exchange buffering is used with blocked records because of the usage of longer channel programs.

The following examples illustrate the control of exchange buffers and the corresponding processing modes that can be used. The buffer pools may have been constructed in any way previously described.

**Exchange Buffering—GET-substitute, PUT-substitute:** The GET macro instruction (step A, Figure 37) specifies the address of a work area. The work area address is exchanged with the address of the next input record returned in register 1. After processing, the address of a record can be passed to the PUT macro instruction. If you issue a PUT macro, it need not be for the record just read.

The PUT macro instruction (step B, Figure 37) specifies the address of the output record. The output record address is exchanged for the address of the next output buffer available for use as a work area. The work area address, returned in register 1, is passed to the GET macro instruction (step C, Figure 37) in register 0.

Notice that as the areas are exchanged there is no movement of data. Output records are contained in the original input area and vice versa, but are logically associated with the correct data set.

**Exchange Buffering—GET-locate, PUTX-output:** The GET macro instruction (step A, Figure 38) locates the address of the next input record. The address is returned in register 1. The record must be processed in the buffer segment before the PUTX macro instruction (step B, Figure 38) is issued. The PUTX macro instruction specifies the address of both the input and output data control block. The two buffer segments are exchanged without any movement of data. The GET macro instruction (step C, Figure 38) locates the next record to be processed

```
                          ...
                          LA    0,WORKAREA
              NEXTREC GET INDCB,(0)
                          LR    0,1
                          ...
                          PUT   OUTDCB,(0)
                          LR    0,1
                          B     NEXTREC
              WORKAREA DS CL50
              INDCB     DCB MACRF=(GT),---
              OUTDCB    DCB MACRF=(PT),---
                          ...
```
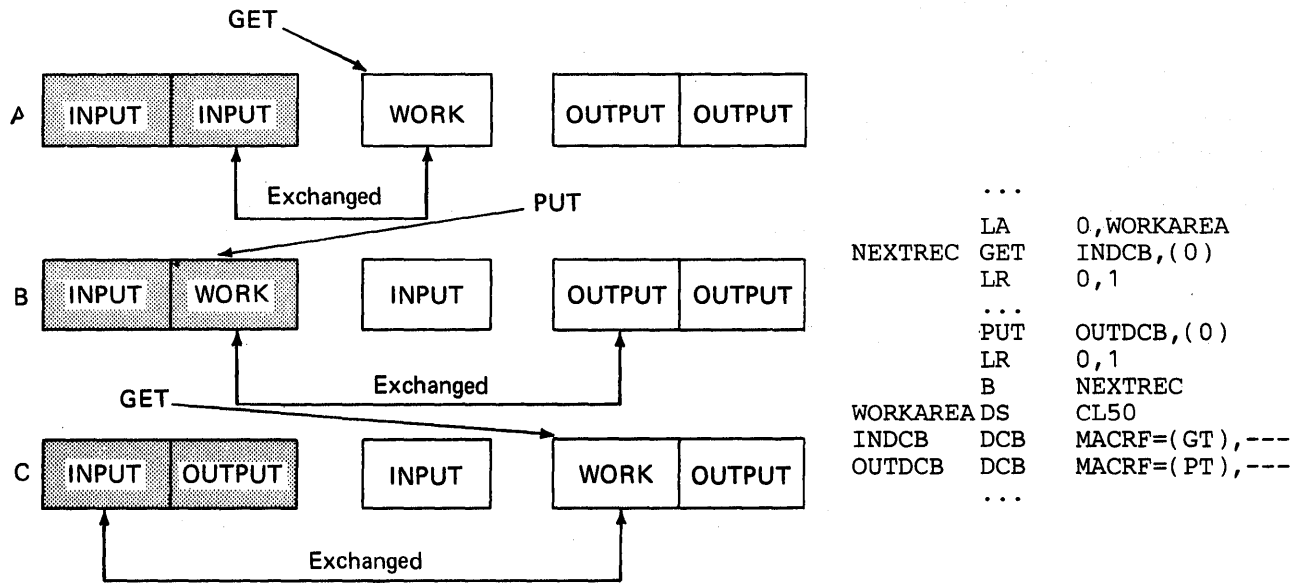
Figure 37. Exchange Buffering with MACRF=GT and MACRF=PT

Notice that the DCB macro instruction for the output data set specifies move mode; this is required.

**Exchange Buffering—GET-locate, PUT-substitute:** The GET macro instruction (step A, Figure 39) locates the next input record. Its address is returned in register 1. You must then move the record to a work area. You can process the record either before or after the move.

The PUT macro instruction (step B, Figure 39) specifies the address of the work area containing the next output record. The system returns the address of the next output



```
              NEXTREC GET  INDCB
                           ...
                      PUTX OUTDCB,INDCB
                      B    NEXTREC
              INDCB    DCB MACRF=(GL),---
              OUTDCB   DCB MACRF=(PM),---
                           ...
```
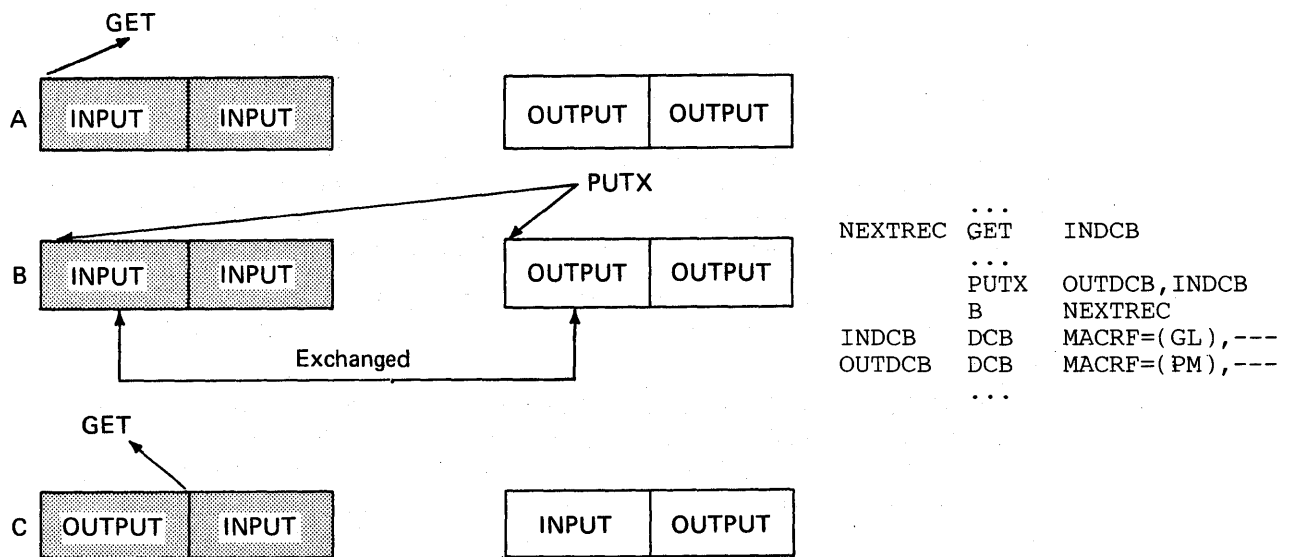
Figure 38. Exchange Buffering with MACRF=GL and MACRF=PM
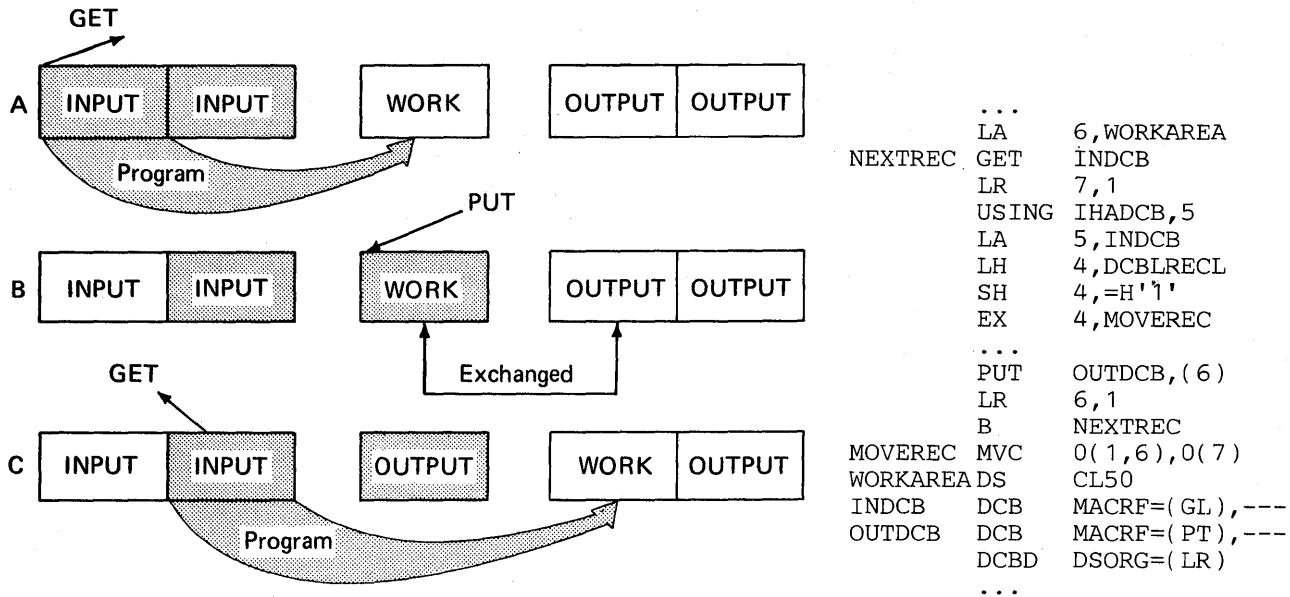
Figure 39. Exchange Buffering with MACRF=GL and MACRF=PT

buffer available for use as a work area in register 1. That address is passed to the move (MVC) instruction in register 6. The GET macro instruction (step C, Figure 39) locates the next input record. You must then move the record to the new work area. Notice that the previous work area becomes part of the output buffer (step C).

Note that if records other than format-F records are being moved, the length attribute of the MVC instruction must be changed as shown by the code beginning with the USING statement in Figure 39. If the record is more than 256 bytes long, you can code a move routine or use an MVCL instruction to process the complete record.

**Buffering Techniques and GET/PUT Processing Modes:** As you can see from the previous examples, the most efficient code is achieved by use of automatic buffer pool construction, and GET-locate and PUTX-output with either simple or exchange buffering. Figure 40 summarizes the combinations of buffering techniques and processing modes that can be used. Notice, for example, that if you use PUT-locate and GET-substitute, you must provide a work area and you must also move the record from the work area to the output buffer.

## RELSE—Release an Input Buffer

When using the queued access technique to process a sequential or indexed sequential data set, you can direct the system to ignore the remaining records in the input buffer being processed. The next GET macro instruction retrieves a record from another buffer. If format-V spanned records are being used, the next logical record obtained may begin on any segment in any subsequent block.

If you are using move mode, the buffer is made available for refilling as soon as the RELSE macro instruction is issued. When you are using locate mode, the system does not refill the buffer until the next GET macro instruction is issued. If a PUTX macro instruction has been used, the block is written before the buffer is refilled.

| Output Buffering: ➤ | Simple | Exchange | Simple | Exchange | Simple | | | | | | Simple | | Exchange | | Simple | | Exchange | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Input Buffering: ➤ Simple / Actions ↓ | GET-move, PUT-locate | GET-move, Put-move | GET-move, PUT-move | GET-move PUT-substitute | GET-locate, PUT-locate | GET-locate, PUT-move | GET-locate, PUT-move | GET-locate, PUT-substitute | GET-locate (logical record), PUT-locate | Input Buffering: Exchange | GET-locate, PUT-locate | GET-locate, PUT-move | GET-locate, PUT-move | GET-locate, PUT-substitute | GET-substitute, PUT-locate | GET-substitute, PUT-move | GET-substitute, PUT-move | GET-substitute, PUT-substitute |
| Program must move record | | | | | X | | | X | X | | X | | X | X | X | | | |
| System moves record | X | X | X | X | | X | X | | | | | X | | | | X | X | |
| System moves record segment | | | | | | | | | X | | | | | | | | | |
| Record is not moved | | | | | | | | | | | | | X | | | | | X |
| Work area required | | X | X | X | | | | X | | | | | | X | X | X | X | X |
| PUTX - output can be used | | | | | | X | X | | | | | X | X | | | | | |

Figure 40. Buffering Technique and GET/PUT Processing Modes

## TRUNC—Truncate an Output Buffer

When using the queued access technique to process a sequential data set, you can direct the system to write a short block. The first record in the next buffer is the next record processed by a PUT-output or PUTX-output mode.

If the locate mode is being used, the system assumes that a record has been placed in the buffer segment pointed to by the last PUT macro instruction.

The last block of a data set is truncated by the Close routine. Note that a data set containing format-F records with truncated blocks cannot be read from direct-access storage as efficiently as a standard format-F data set.

## GETBUF—Get a Buffer from a Pool

The GETBUF macro instruction can be used with the basic access technique to request a buffer from a buffer pool constructed by the BUILD, GETPOOL, or OPEN macro instruction. The address of the buffer is returned by the system in a register you specify when you issue the macro instruction. If no buffer is available, the register contains 0 instead of an address.

## FREEBUF—Return a Buffer to a Pool

The FREEBUF macro instruction is used with the basic access technique to return a buffer to the buffer pool from which it was obtained by a GETBUF macro instruction. Although the buffers need not be returned in the order in which they were obtained, they must be returned when they are no longer needed.

## FREEDBUF—Return a Dynamic Buffer to a Pool

Any buffer obtained through the dynamic buffer option must be released before it can be used again. When you are processing a direct data set, if you do not update the block in the buffer and thus free the buffer when the block is written, you must use the FREEDBUF macro instruction. If there is an uncorrectable input/output error, the control program releases the buffer. When you are processing an indexed sequential data set, if you do not update the block in the buffer or if there is an uncorrectable input error, the control program releases the buffer when the next READ macro instruction is issued on the same DECB, unless you use the FREEDBUF macro instruction.

To effect the release, you must specify the address of the DECB that was used when the block was read using the dynamic buffering option, as well as the address of the DCB associated with the data set being processed.

# Processing a Sequential Data Set

Data sets residing on all volumes other than direct-access volumes must be processed sequentially. In addition, a data set residing on a direct-access volume, regardless of organization, can be processed sequentially. This includes data sets created using ISAM or a similar access method. Since the entire data set (prime, index, and overflow areas) will be processed, care should be taken to determine the type of records being processed. This feature of the operating system allows you to write your program with little regard for the type of device to be used when the program is executed, except for restrictions on the use of certain device-dependent macro instructions and processing options.

Either the queued or the basic access technique may be used to store and retrieve the records of a sequential data set. Additionally, a technique called *chained scheduling* can be used to accelerate the input/output operations required for a sequential data set.

## *Data Format—Device Type Considerations*

Before execution of your program, you must supply the operating system with both the record format (RECFM) and device-dependent (DEVD) information in a DCB macro instruction, a DD statement, or a data set label. The DCB subparameters for the DD statement differ slightly from those described here. A complete description of the DD statement and a glossary of DCB subparameters are contained in *OS/VS1 JCL Reference*.

The record format (RECFM) parameter of the DCB macro instruction specifies the characteristics of the records in the data set as fixed-length (RECFM=F), variable-length (RECFM=V or D), or undefined-length (RECFM=U). Fixed-length blocked records (RECFM=FB) can be specified as standard (RECFM=FBS), which means there are no truncated (short) blocks or unfilled tracks within the data set, with the possible exception of the last block or track. Data sets with a fixed-length, standard format were described previously under "Fixed-Length Records, Standard Format."

As an optional feature, a control character can be contained in each record. This control character will be recognized and processed if the data set is printed or punched. The control characters are transmitted on both tapes and direct-access volumes. The presence of a control character is indicated by M or A in the RECFM field of the data control

block. M denotes machine code; A denotes American National Standards Institute (ANSI) code. If either M or A is specified, the character must be present in every record; the printer space (PRTSP) or stacker select (STACK) field of the DCB is ignored. The optional control character must be in the first byte of format-F and format-U records and in the fifth byte of format-V records and format-D records where BUFOFF=L. Control character codes are listed in "Appendix B: Control Characters." The device-dependent (DEVD) parameter of the DCB macro instruction specifies the type of device on which the data set's volume resides:

TA magnetic tape
PT paper tape reader
PR printer
PC card punch
RD card reader
DA direct-access device or
   Mass Storage System (MSS) virtual volumes

## Magnetic Tape (TA)

Format-F, V, D, and U records are acceptable for magnetic tape. Format-V records are not acceptable on 7-track tape if the data conversion feature is not available. ASCII records are not acceptable on 7-track tape.

When you create a tape data set with variable-length record format (V or D), the control program pads any data block shorter than 18 bytes. For format-V records, it pads to the right with binary zeros so that the data block length equals 18 bytes. For format-D (ASCII) records, the padding consists of ASCII circumflex (>) characters which are equivalent to X'5E's.

Note that there is no minimum requirement for blocksize; however, if a data check occurs on a magnetic-tape device, any record shorter than 12 bytes in a read operation or 18 bytes in a write operation will be treated as a noise record and lost. No check for noise will be made unless a data check occurs.

Tape density (DEN) specifies the recording density in bits per inch per track, as shown in Figure 41. If this information is not supplied, the highest applicable density is assumed.

---

**Recording Density**

| DEN | 7-Track Tape | 9-Track Tape |
|-----|--------------|--------------|
| 0 | 200 | — |
| 1 | 556 | — |
| 2 | 800 | 800 (NRZI) |
| 3 | — | 1600 (PE) |
| 4 | — | 6250 (GCR) |

NRZI is for non-return-to-zero-inverted mode
PE is for phase encoded mode
GCR is for group coded recording mode

**Note:** Specifying DEN=0 for a 7-track 3420 tape attached to a 3803-1 will result in 556 bits per inch recording, but corresponding messages and tape labels will indicate 200 bits per inch recording density.

Figure 41 Tape Density (DEN) Values

---

The track recording technique (TRTCH) for 7-track tape can be specified as:

C Data conversion is to be used. Data conversion makes it possible to write 8 binary bits of data on 7 tracks. Otherwise, only 6 bits of an 8-bit byte are recorded. The length field of format-V records contains binary data and is not recorded correctly without data conversion.

E Even parity is to be used; if E is omitted, odd parity is assumed.

T    BCDIC to EBCDIC translation is required.

**Paper-Tape Reader (PT)**

The paper-tape reader accepts format-F and format-U records. If you use QSAM, you should not specify the records as blocked. Each format-U record is followed by an end-of-record character. Data read from paper tape may optionally be converted into the System/370 internal representation of one of six standard paper-tape codes. Any character found to have a parity error will not be converted when the record is transferred into the input area. Characters deleted in the conversion process are not counted in determining the block size.

The following symbols indicate the code in which the data was punched. If this information is omitted, I is assumed.

    I    IBM BCD perforated tape and transmission code (8 tracks)
    F    Friden (8 tracks)
    B    Burroughs (7 tracks)
    C    National Cash Register (8 tracks)
    A    ASCII (8 tracks)
    T    Teletype[1] (5 tracks)
    N    No conversion

Note that when you are using QSAM, the processing mode must be move mode.

**Card Reader and Punch (RD/PC)**

Format-F and U records are acceptable to both the reader and punch; format-V records are acceptable to the punch only. The device control character, if specified in the RECFM parameter, is used to select the stacker; it is not punched. The first 4 bytes (record descriptor word or segment descriptor word) of format-V records or record segments are not punched. For format-V records, at least 1 byte of data must follow the record or segment descriptor word or the carriage control character.

Each punched card corresponds to one physical record. Therefore, you should restrict the maximum record size to 80 (EBCDIC mode) or 160 (column binary mode) data bytes. When mode (C) is used for the card punch, BLKSIZE must be 160 unless you are using PUT. Then you can specify BLKSIZE as 160 or a multiple of 160, and the system handles this as described earlier under "PUT—Write a Record" in the section "Queued Access Techniques." You can specify the read/punch mode of operation (MODE) parameter as either card image (column binary) mode (C) or EBCDIC mode (E). If this information is omitted, E is assumed. The stacker selection parameter (STACK) can be specified as either 1 or 2 to indicate which bin is to receive the card. If it is not specified, 1 is assumed.

For all QSAM, RECFM=FB, card punch data sets, the block size in the DCB will be adjusted by the system to equal the logical record length. This data set will be treated as RECFM=F. If the system builds the buffers for this data set, the buffer length will be determined by the BUFL parameter. If the BUFL parameter was not specified, the adjusted block size is used for the buffer length.

If the DCB is to be reused with a block size larger than the logical record length, you must reset DCBBLKSI in the DCB and insure that the buffers are large enough to contain the largest block size expected. You may insure the buffer size by specifying the BUFL parameter before the first time the data set is opened or by issuing the FREEPOOL macro instruction after each CLOSE macro so the system will build a new buffer pool of the correct size each time the data set is opened.

---

[1] Trademark of the Teletype Corporation

Note that when QSAM is used, punch error correction on the IBM 2540 Card Read Punch is automatically performed only for data sets using three or more buffers without the chained scheduling feature.

The 3525 Card Punch accepts only format-F records for print data sets and for associated data sets. Other record formats are allowed for the read data set, the punch data set, and the interpret punch data set. See *OS and OS/VS Programming Support for the IBM 3505 Card Reader and IBM 3525 Card Punch* for more information on programming for the 3525 Card Punch.

## Printer (PR)

Records of format-F, V, and U are acceptable to the printer. The first 4 bytes (record descriptor word or segment descriptor word) of format-V records or record segments are not printed. For format-V records, at least 1 byte of data must follow the record or segment descriptor word or the carriage control character. The carriage control character, if specified in the RECFM parameter, is not printed. The system does not position the printer to channel 1 for the first record unless specified by a carriage control character.

Because each line of print corresponds to one record, the record length should not exceed the length of one line on the printer. For variable-length spanned records, each line corresponds to one record segment, and blocksize should not exceed the length of one line on the printer.

If carriage control characters are not specified, you can indicate printer spacing (PRTSP) as 0, 1, 2, or 3. If it is not specified, 1 is assumed.

For all QSAM, RECFM=FB, printer data sets, the block size in the DCB will be adjusted by the system to equal the logical record length. This data set will be treated as RECFM=F. If the system builds the buffers for this data set, the buffer length will be determined by the BUFL parameter. If the BUFL parameter was not specified, the adjusted block size is used for the buffer length.

If the DCB is to be reused with a block size larger than the logical record length, you must reset DCBBLKSI in the DCB and insure that the buffers are large enough to contain the largest block size expected. You may insure the buffer size by specifying the BUFL parameter before the first time the data set is opened or by issuing the FREEPOOL macro instruction after each CLOSE macro so the system will build a new buffer pool of the correct size each time the data set is opened.

## Direct-Access Device (DA)

Direct-access devices accept records of format-F, V, or U. If the records are to be read or written with keys, the key length (KEYLEN) must be specified. In addition, the operating system has a standard track format for all direct access volumes. Each track contains data information as well as certain control information such as:

- The address of the track
- The address of each record
- The length of each record
- Gaps between areas

A complete description of track format is contained in the section "Direct-Access Device Characteristics." Your only concern in creating a sequential data set is to allow for an 8-byte track descriptor record (capacity record or R0) when requesting space on a direct-access volume. In addition, device overhead, which varies with the device, must be allocated for each block on the track.

## Device Control

The operating system provides you with six macro instructions for controlling input/output devices. Each is, to varying degrees, device-dependent. Therefore, you must exercise some care if you wish to achieve device independence.

When you use the queued access technique, only unit record equipment can be controlled directly. When using the basic access technique, limited device independence can be achieved between magnetic-tape and direct-access devices. You must check all read or write operations before issuing a device control macro instruction.

### CNTRL—Control an I/O Device

The CNTRL macro instruction performs these device-dependent control functions:

- Card reader stacker selection (SS)

- Printer line spacing (SP)

- Printer carriage control (SK)

- Magnetic-tape backspace (BSR) over a specified number of blocks

- Magnetic-tape backspace (BSM) past a tapemark and forward space over the tapemark

- Magnetic-tape forward space (FSR) over a specified number of blocks

- Magnetic-tape forward space (FSM) past a tapemark and a backspace over the tapemark

Backspacing moves the tape toward the load point; forward spacing moves the tape away from the load point.

Note that the CNTRL macro instruction cannot be used with an input data set containing variable-length records on the card reader.

You can use the CNTRL macro instruction to position DOS tapes that contain embedded DOS checkpoint records if you specify OPTCD=H in the DCB parameter field of the DD statement. The CNTRL macro instruction cannot be used to backspace DOS 7-track tapes that are written in data convert mode and contain embedded checkpoint records.

### PRTOV—Test for Printer Overflow

The PRTOV macro instruction tests for channel 9 or 12 of the printer carriage control tape or the forms control buffer (FCB). An overflow condition causes either an automatic skip to channel 1 or, if specified, transfer of control to your routine for overflow processing. If you specify an overflow exit routine, set DCBIFLGS to X'00' before issuing another PRTOV.

If the data set specified in the DCB is not for a printer, no action is taken.

### SETPRT—Printer Setup

The SETPRT macro instruction is·used to initially set or dynamically change the specifications of the 3800 Printing Subsystem. For additional information on how to use the SETPRT macro with the 3800 printer, see *IBM 3800 Printing Subsystem Programmer's Guide*.

For printers that have a universal character set (UCS) buffer or a forms control buffer (FCB), the SETPRT macro instruction is used to fetch UCS and FCB images from the image library (SYS1.IMAGELIB) and load them into their respective buffers. Note that FCB images for the 3203 or 3211 and 3800 are not compatible. The universal character

sets for the 1403, or 3203 and the character arrangement tables for the 3800 are also incompatible.

The SETPRT macro allows you to request the operator to verify loading of the buffer. For the 1403, 3203, and 3211 printers, the SETPRT macro allows you to specify the printing of lowercase EBCDIC characters in uppercase when no uppercase/lowercase print chain or train is available.

For a printer that has no carriage control tape, you can use the SETPRT macro instruction to load the FCB, to request operator verification of buffer loading, and to allow the operator to align the paper in the printer.

The FCB contents can be fetched from the system library or defined in your program through the exit list of the DCB macro instruction, as discussed under "Exit List (EXLST)."

When issued, the SETPRT macro instruction loads a special UCS buffer from the system library. The library contains images of standard IBM character sets and of your own special character sets. The operator can be requested to verify the loaded image after mounting the appropriate print chain or train.

The SETPRT macro instruction can be used to block or unblock printer data checks. When data checks are blocked, unprintable characters are treated as blanks and do not cause an error condition.

If the specified UCS or FCB image is not found in the image library (or DCB exit list for an FCB image), the operator is requested to specify a different one (message IEC127D is issued).

## BSP—Backspace a Magnetic Tape or Direct-Access Volume

The BSP macro instruction backspaces one block on the magnetic tape or direct-access volume being processed. The block can then be reread or rewritten. An attempt to rewrite the block destroys the contents of the remainder of the tape or track.

The direction of movement is toward the load point or beginning of the extent. You may not use the BSP macro instruction if the track overflow option was specified or if the CNTRL, NOTE, or POINT macro instruction is used. The BSP macro instruction should be used only when other device control macro instructions could not be used for backspacing.

Any attempt to backspace across a file mark will result in a return code of X'04' and your tape or direct-access volume will be positioned after the file mark. This means you cannot issue a successful backspace command once your EODAD routine is entered unless you first reposition the tape or direct-access volume into your data set. (CLOSE TYPE=T can get you positioned at the end of your data set.)

You can use the BSP macro instruction to backspace DOS tapes containing embedded DOS checkpoint records. If you use this means of backspacing, you must test for and bypass the embedded checkpoint records. You cannot use the BSP macro instruction for DOS 7-track tapes written in translate mode.

## NOTE—Return the Relative Address of a Block

The NOTE macro instruction requests the relative address of the block just read or written. In a multivolume data set, the address is relative to the beginning of the volume currently being processed.

The address provided by the operating system is returned in register 1. The address is in the form of a 4-byte relative block address for magnetic tape; for a direct-access device, it is a 4-byte relative track address. The amount of unused space available on the track of the direct-access device is returned in register 0.

## POINT—Position to a Block

The POINT macro instruction causes repositioning of a magnetic tape or direct-access volume to a specified block. The next read or write operation begins at this block. In a multivolume data set, you must ensure that the volume referred to is the volume currently being processed. If a write operation follows the POINT macro instruction, all of the track following the write operation is erased unless the data set is opened for UPDAT. POINT is not meant to be used before a WRITE macro instruction when a data set is opened for UPDAT. You can use the POINT macro instruction to position DOS tapes that contain embedded checkpoint records if you specify OPTCD=H in the DCB parameter field of the DD statement. The POINT macro instruction cannot be used to backspace DOS 7-track tapes that are written in data convert mode and contain embedded checkpoint records.

When using the POINT macro for a direct-access device that is opened for OUTPUT, OUTIN, or INOUT, and the record format is not standard, the number of blocks per track may vary slightly.

## *Device Independence*

The ability to request input/output operations without regard for the physical characteristics of the I/O devices makes it possible for you to write one program that will fulfill a variety of needs. Device independence may be useful for:

- Accepting data from a number of recording devices, such as a disk pack, 7- or 9-track magnetic tape, or unit-record equipment. This situation could arise when several types of data-acquisition devices are feeding a centralized complex.

- Observing constraints imposed by the availability of input/output devices (for example, when devices on order have not been installed).

- Assembling, testing, and debugging on one System/370 configuration and processing on a different configuration. For example, a 2314 drive can be used as a substitute for several magnetic-tape units.

Device independence is not difficult to achieve, but requires some planning and forethought. There are two areas of planning necessary to achieve device independence—system generation considerations and programming considerations.

### System Generation Considerations

You can provide for device independence when the system is generated by generating a system that not only meets the current input/output configuration requirements but includes anticipated device attachments. Creating such a system entails looking ahead at expected delivery of input/output devices and, during system generation, constructing the necessary control blocks and tables. Thus, when the devices are delivered, they need only be physically attached. The operating system recognizes the devices without modification. However, until the devices are physically connected, the operator must designate them as being offline, using the VARY command or using automatic device status initialization. For information on automatic device status initialization, see *OS/VS1 System Generation Reference.*

When new device attachments cannot be fully anticipated, you can add new devices by performing an I/O device generation. This is a limited type of system generation that enables you to change your I/O configuration without regenerating other parts of the system.

System generation techniques for effecting a smooth transition to new input/output devices do not include addition of new device types. When support for new devices is

provided, a new system must be generated. A complete description of system generation techniques is contained in *OS/VS1 System Generation Reference.*

## Programming Considerations

Each of three data set organizations—partitioned, indexed sequential, and direct—requires the use of a direct-access device. Device independence is meaningful, then, only for a sequentially organized data set, that is, a data set where one block of data follows another, thus allowing input or output to be on a magnetic tape drive, a direct-access device, a card read/punch, a printer, or a spooled data set.

Your program will be device-independent if you do two things:

- Omit all device-dependent macro instructions and macro instruction parameters from your program.

- Defer specifying any required device-dependent parameters until the program is ready for execution. That is, supply the parameters on your data definition (DD) statement or during the open exit routine.

In examining the following list of macro instructions, consider only the logical layout of your data record without regard for the type of device used. Also, consider that any reference to a direct-access volume is to be treated like a reference to magnetic tape, that is, you must create a new data set rather than attempt to update.

OPEN
| Specify INPUT, OUTPUT, INOUT, OUTIN, OUTINX, or EXTEND. The parameters RDBACK and UPDAT are device-dependent and cause an abnormal termination if directed to a device of the wrong type.

READ
Specify forward reading (SF) only.

WRITE
Specify forward writing (SF) only; use only to create new records.

PUTX
Use only output mode.

NOTE/POINT
These macros are valid for both magnetic-tape and direct-access volumes.

BSP
This macro is valid for magnetic-tape or direct-access volumes. However, its use would be an attempt to perform device-dependent action.

CNTRL/PRTOV
These macros are device-dependent.

### DCB Subparameters

MACRF
Specify R/W or G/P. Processing mode can also be indicated.

DEVD
Specify DA if any direct-access device may be used. Magnetic-tape and unit-record equipment DCBs will fit in the area provided during assembly. Specify unit-record devices only if you expect never to change to tape or direct-access devices. Key length (KEYLEN) can be specified on the DD statement if necessary.

RECFM, LRECL, BLKSIZE
   These can be specified in the DD statement. However, you must consider maximum
   record size for specific devices, and track overflow cannot be specified unless
   supported.

DSORG
   Specify sequential organization (PS or PSU).

OPTCD
   This subparameter is device-dependent; specify it in the DD statement.

SYNAD
   Any device-dependent error checking is automatic. Generalize your routine so that no
   device-dependent information is required.

## Chained Scheduling for I/O Operations

To accelerate the input/output operations required for a data set, the operating system
provides a technique called *chained scheduling*. When requested, the system bypasses
the normal I/O routines and dynamically chains several input/output operations
together. A series of separate read or write operations, functioning with chained
scheduling, is issued to the computing system as one continuous operation. In a
nonpageable partition or address space, the program-controlled interruption (PCI) flag in
the CCWs is used for synchronization of the I/O operations.

The I/O performance is improved by reduction in both the CPU time and the channel
start/stop time required to transfer data within virtual storage. Some factors that affect
performance improvement are:

- Address space type (real or virtual)

- BUFNO for QSAM

- The number of overlapped requests for BSAM

- Other activity on the CPU and channel

The effects of rotational delay are also reduced since several successive blocks, requested
separately, can be retrieved in a single rotation. Chained scheduling can be used only
with simple buffering. Each data set for which chained scheduling is specified must be
assigned at least two and preferably three buffers with QSAM, or must have a value of at
least two and preferably three for NCP with BSAM or BPAM.

Chained scheduling will not be used where it is not allowed.

A request for chained scheduling will be ignored and normal scheduling used if any of
the following are encountered when the data set is opened:

- BDAM CREATE, that is, MACRF=(WL)

- Track overflow

- UPDAT in the operand of the OPEN macro instruction

- Exchange buffering

- CNTRL macro instruction to be used

- Device type is paper tape reader

- Bypassing of embedded DOS checkpoint records on tape input data sets

- Spooled data sets (SYSIN or SYSOUT)

- A print data set or any associated data set for the 3525 Card Punch. (See *OS and OS/VS Programming Support for the IBM 3505 Card Reader and IBM 3525 Card Punch* for more information on programming for the 3525.)

The number of channel program segments that can be chained is limited to the value specified in the NCP operand of BSAM and BPAM DCBs, and to the value specified in the BUFNO operand of QSAM DCBs.

When chained scheduling is being used, the automatic skip feature of the PRTOV macro instruction for the printer will not function. Format control must be achieved by ANSI or machine control characters. (Control characters are discussed in more detail in Part 1 under "Control Character," in Part 2 under "Data Format—Device Type Considerations," and in "Appendix B: Control Characters.") When you read undefined-length records with QSAM, the DCBLRECL field is equal to the BLKSIZE field, not the actual record length. The entire block is moved to your work area in the move mode. When you are using QSAM under chained scheduling to read variable-length, blocked, ASCII tape records (format-DB), you must code BUFOFF=L in the DCB for that data set.

Note also that if you are using BSAM with the chained scheduling option to read format-DB records and have coded a value for the BUFOFF operand other than BUFOFF=L, the input buffers will be converted from ASCII to EBCDIC as usual, but the record length returned to the DCBLRECL field will equal the block size, not the actual length of the record read in; the record descriptor word (RDW), if present, will not have been converted from ASCII to binary.

When chained scheduling is used on the 2540 Card Read Punch, error recovery procedures are not performed.

Chained scheduling is most valuable for programs that require extensive input and output operations. Because a data set using chained scheduling may monopolize available time on a channel, separate channels should be assigned, if possible, when more than one data set is to be processed.

## Search Direct for Input Operations

To accelerate the input operations required for a data set on DASD, the operating system provides a technique called *search direct*. Search direct reads in the requested record and the count field of the second record. This allows the operation to get the next record directly, along with the count field of the following record. Search direct can be used with all record formats except format-UT, format-FBT, format-FS, format-FBS, and spanned. You request search direct by coding OPTCD=Z in the DCB macro instruction. For FS and FBS records, the access method routines always use a form of search-direct processing. Search direct cannot be used under the following conditions:

- In conjunction with the NOTE and POINT macro instructions
- When you specify the UPDAT option of the OPEN macro instruction
- For partitioned data sets

## Creating a Sequential Data Set

As discussed earlier, a processing program should be developed using, as much as possible, factors that are constant, with variable factors specified at execution. For that reason, the following examples are generalized as much as possible. They are neither exhaustive nor intended as complete examples. Rather, they are presented as introductory sequences.

In creating a sequential data set on a magnetic tape or direct-access device, you must do the following:

- Code DSORG=PS or PSU in the DCB macro instruction

- Code a DD statement to describe the data set (See *OS/VS1 JCL.*)

- Process the data set with an OPEN macro instruction (data set is opened for output or OUTIN), a series of PUT or WRITE and CHECK macros, and then a CLOSE macro

**Tape-to-Print, Move Mode—Simple Buffering:** In Figure 42, the GET-move and PUT-move require two movements of the data records. If the record length (LRECL) does not change in processing, only one move is necessary; you can process the record in the input buffer segment. A GET-locate can be used to provide a pointer to the current segment.

**Tape-to-Print, Locate Mode—Simple Buffering:** This example (Figure 43) is similar to that in Figure 42. However, since there is no change in the record length, the records can be processed in the input buffer. Only one move of each data record is required.

**Tape-to-Print, Substitute Mode—Exchange Buffering:** Although the initial problem is the same, the solution described in Figure 44 takes advantage of exchange buffering, which eliminates the need to move the data record, and makes direct reference to individual fields within a record through the use of a dummy control section (DSECT). The use of the DSECT allows symbolic reference to be made for storage-to-storage operations; therefore, the length attributes need not be explicitly stated.

```
          ...
          OPEN       ( INDATA, ,OUTDATA, ( OUTPUT ) )
NEXTREC   GET        INDATA,WORKAREA                    Move mode
          AP         NUMBER,=P'1'
          UNPK       COUNT,NUMBER                       Record count adds 6
          PUT        OUTDATA,COUNT                      bytes to each record
          B          NEXTREC
TAPERROR  SYNADAF    ACSMETH=QSAM                       Control program returns message
          LA         0,68( 0,1 )                        address in register 1.
          ST         14,SAVE14                          SYNAD routine prints part of
          PUT        OUTDATA,( 0 )                      the message ( beginning with
          SYNADRLS                                      the unit number ) as a 56-byte
          L          14,SAVE14                          fixed-length record. It then
          RETURN                                        returns to the control
ENDJOB    CLOSE      ( INDATA, ,OUTDATA )               program.
          ...
WORKAREA  DS         CL50
COUNT     DS         CL6
NUMBER    DC         PL4'0'
SAVE14    DS         F
INDATA    DCB        DDNAME=INPUTDD,DSORG=PS,MACRF=( GM ),EROPT=ACC,          C
                     SYNAD=TAPERROR,EODAD=ENDJOB
OUTDATA   DCB        DDNAME=OUTPUTDD,DSORG=PS,MACRF=( PM ),EROPT=ACC
          ...
```

Figure 42. Creating a Sequential Data Set—Move Mode, Simple Buffering

```
              . . .
              OPEN        ( INDATA,,OUTDATA,(OUTPUT),ERRORDCB,(OUTPUT))
NEXTREC       GET         INDATA                          Locate mode
              LR          2,1                             Save pointer
              AP          NUMBER,=P'1'
              UNPK        0(6,2),NUMBER                   Process in input area
              PUT         OUTDATA                         Locate mode
              MVC         0(50,1),0(2)                    Move record to output buffer
              B           NEXTREC
TAPERROR      SYNADAF     ACSMETH=QSAM                    Message address in register 1
              ST          2,SAVE2                         Save register 2 contents
              L           2,8(0,1)                        Load pointer to input buffer
              MVC         8(70,1),50(1)                   Shift nonblank message fields
              MVI         78(1),C' '                      Blank end of message
              MVC         79(49,1),78(1)
              ST          2,128(1)                        Save address for debugging
              CH          0,=H'4'                         Test SYNADAF return code
              BE          MOVERCD                         Branch if data read
              BL          PRINTIT                         Branch if data not read
              CLI         128(1),C' '                     See if data read anyway
              BE·         PRINTIT                         Branch if definitely no data
MOVERCD       MVC         78(50,1),0(2)                   Add input record to message
PRINTIT       LA          0,4(1)                          Load address of message
              LR          2,14                            Save return address
              PUT         ERRORDCB,(0)                    Print message (move mode)
              SYNADRLS                                    Release message and save area
              LR          14,2                            Restore return address
              L           2,SAVE2                         Restore register 2 contents
              RETURN                                      Return to control program
ENDJOB        CLOSE       ( INDATA,,OUTDATA,,ERRORDCB)
              . . .
NUMBER        DC          PL4'0'
INDATA        DCB         DDNAME=INPUTDD,DSORG=PS,MACRF=(GL),EROPT=ACC,          C
                          SYNAD=TAPERROR,EODAD=ENDJOB
OUTDATA       DCB         DDNAME=OUTPUTDD,DSORG=PS,MACRF=(PL)
ERRORDCB      DCB         DDNAME=SYSOUTDD,DSORG=PS,MACRF=(PM),RECFM=V,           C
                          BLKSIZE=128,LRECL=124
SAVE2         DS          F
              . . .
```

Figure 43. Creating a Sequential Data Set—Locate Mode, Simple Buffering

## Retrieving a Sequential Data Set

In retrieving a sequential data set on a magnetic tape or direct-access device, you must do the following:

- Code DSORG=PS or PSU in the DCB macro instruction

- Tell the system where your data set is located (by coding a DD statement; see *OS/VS1 JCL Reference.*

- Process the data set with an OPEN macro instruction (data set is opened for input, INOUT, RDBACK, or UPDAT), a series of GET or READ macros and then a CLOSE macro.

## Updating a Sequential Data Set

When you update in place, you read records, process them, and write them back to their original positions without destroying the remaining records on the track. The following rules apply:

- You must specify the update option (UPDAT) in the OPEN macro instruction. To perform the update, you can use only the READ, WRITE, CHECK, NOTE, POINT, GET, and PUTX macro instructions.

```
            . . .
            OPEN         ( INDATA, , OUTDATA, ( OUTPUT ) , ERRORDCB, ( OUTPUT ) )
            LA           0,GIVEAWAY                      Set up for first buffer
NEXTREC     GET          INDATA,( 0 )                    Substitute mode
            LR           2,1                             Pointer to next record
            USING        RECORD,2                        Establish address of DSECT
            AP           NUMBER,=P'1'
            UNPK         COUNT,NUMBER
            PUT          OUTDATA,RECORD                  Substitute mode
            LR           0,1                             Exchange work area
            B            NEXTREC
TAPERROR    SYNADAF      ACSMETH=QSAM                    SYNAD routine is same
            . . .                       as in previous example
ENDJOB      CLOSE        ( INDATA, , OUTDATA, , ERRORDCE )
            . . .
            DS           0D
GIVEAWAY    DS           CL50
NUMBER      DC           PL4'-'
INDATA      DCB          DDNAME=INPUTDD,DSORG=PS,MACRF=( GT ),BFTEK=E,BFALN=D,          C
                         EROPT=ACC,SYNAD=TAPERROR,EODAD=ENDJOB
OUTDATA     DCB          DDNAME=OUTPUTDD,DSORG=PS,MACRF=( PT ),BFTEK=E,BFALN=D,         C
                         EROPT=ACC

            . . .
RECORD      DSECT
COUNT       DS           CL6
RESTOFIT    DS           CL44
            . . .
```

Figure 44. Creating a Sequential Data Set—Substitute Mode, Exchange Buffering

- You cannot use chained scheduling.

- You cannot delete any record or change its length; you cannot add new records.

- The data set must be on a direct-access device.

A record must be retrieved by a READ or GET macro instruction before it can be updated by a WRITE or PUTX macro instruction. A WRITE or PUTX macro instruction does not need to be issued after each READ or GET macro instruction. The READ and WRITE macro instructions must be execute forms that refer to the same DECB; the DECB must be provided by the list forms of the READ or WRITE macro instructions. (The execute and list forms of the READ and WRITE macro instructions are described in *OS/VS1 Data Management Macro Instructions.*)

**Updating With Overlapped Operations:** To overlap input/output and CPU activity, you can start several read or write operations before checking the first for completion. You cannot overlap read with write operations, however, as operations of one type are started or resumed. Note that each concurrent read or write operation requires a separate channel program and a separate DECB. If a single DECB were used for successive read operations, only the last record read could be updated.

In Figure 54, overlap is achieved by having a read or write request outstanding while each record is being processed. Note the use of the execute and list forms of the READ and WRITE macro instructions, identified by the operands MF=E and MF=L.

## Extending a Sequential Data Set

If you want to add records at the end of your data set, you must open the data set for output with DISP=MOD specified in the DD statement, or the EXTEND option specified in the OPEN macro. You can then issue PUT or WRITE macros to the data set.

## Determining the Length of a Record When Using the BSAM READ Macro

When you read a sequential data set, you can determine the length of the record in one of the following four ways, depending upon the record format of the data set:

- For fixed-length, unblocked records, the length of all records is the value in the DCBBLKSI field of the DCB.

- For variable-length records, the block descriptor word in the record contains the length of the record.

- For fixed-length blocked or undefined-length records, the following method can be used to calculate the block length. (This method should not be used when reading track overflow records on a device with the rotational position sensing (RPS) feature or when using chained scheduling on any device. The length of a record cannot be determined when using chained scheduling.) After checking the DECB for the READ request but before issuing any subsequent data management macro instructions that specify the DCB for the READ request, obtain the IOB address from the DECB. The IOB address can be loaded from the location 16 bytes from the start of the DECB.

  Obtain the residual count from the channel status word (CSW) that has been stored in the input/output block (IOB). The residual count is in the halfword 14 bytes from the start of the IOB. Subtract this residual count from the number of data bytes requested to be read by the READ macro instruction. If "S" was coded as the length parameter of the READ macro instruction, the number of bytes requested is the value of DCBBLKSI at the time the READ was issued. If the length was coded in the READ macro instruction, this value is the number of data bytes and it is contained in the halfword 6 bytes from the beginning of the DECB. The result of the subtraction is the length of the block read. See Figure 45.

```
        ...
        OPEN(DCB,(INPUT))
        LA          DCBR,DCB
        USING       IHADCB,DCBR
        ...
        READ        DECB1,SF,DCB,AREA1,'S'
        READ        DECB2,SF,DCB,AREA2,50
        ...
        CHECK       DECB1
        LH          WORK1,DCBBLKSI              Block size at time of READ
        L           WORK2,DECB1+16             IOB address
        SH          WORK1,14(WORK2)            WORK1 has block length
        ...
        CHECK       DECB2
        LH          WORK1,DECB2+6              Length requested
        L           WORK2,DECB2+16            IOB address
        SH          WORK1,14(WORK2)           WORK1 has block length
        ...
        MVC         DCBBLKSI,LENGTH3           Length to be read
        READ        'DECB3,SF,DCB,AREA3
        ...
        CHECK       DECB3
        LH          WORK1,LENGTH3             Block size at time of READ
        L           WORK2,DECB+16             IOB address
        SH          WORK1,14(WORK2)           WORK1 has block length
        ...
DCB     DCB         ...RECFM=U,NCP=2,...
        DCBD
        ...
```

Figure 45. One Method of Determining the Length of the Record When Using BSAM to Read Undefined-Length Records

- For undefined-length records, the LRECL operand should be omitted; the actual length can be supplied dynamically in a READ/WRITE macro instruction. (This method should not be used when reading track overflow records on a device with the rotational position sensing (RPS) feature or when using chained scheduling on any device.) When an undefined-length record is read, the actual length of the record is returned by the system in the DCBLRECL field of the data control block.

### *Writing a Short Block When Using the BSAM WRITE Macro*

When you are writing blocks for a sequential data set, you can change the length of a block if you have fixed-blocked record format. The DCB block size field (DCBBLKSI) can be changed to specify a block size that is shorter than what was originally specified for the data set. The DCBBLKSI field must be changed before issuing the WRITE macro instruction and must be a multiple of the LRECL parameter in the DCB. Any subsequent WRITE macro instructions issued will write records with the new block length until the block size is changed again. The DCB block size field should not be changed to specify a block size that is greater than what was originally specified for the data set.

## Processing a Partitioned Data Set

A partitioned data set can be stored only on a direct-access device. It is divided into sequentially organized *members*, each made up of one or more records (see Figure 46). Each member has a unique name, 1 to 8 characters long, stored in a *directory* that is part of the data set. The records of a given member are stored or retrieved sequentially.

The main advantage of using a partitioned data set is that you can retrieve any individual member once the data set is opened without searching the entire data set. For example, a program library can be stored as a partitioned data set, each member of which is a separate program or subroutine. The individual members can be added or deleted as required. When a member is deleted, the member name is removed from the directory, but the space used by the member cannot be reused until the data set is reorganized.
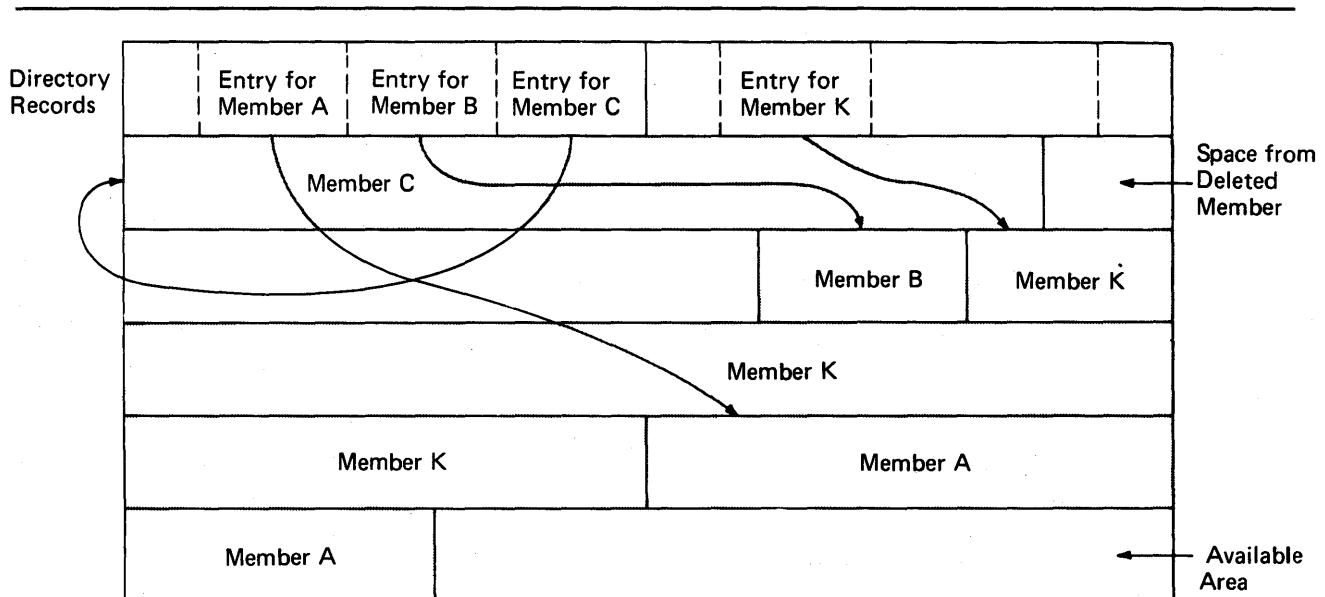


Figure 46. A Partitioned Data Set

The directory, a series of records at the beginning of the data set, contains an entry for each member. Each directory entry contains the member name and the starting location of the member within the data set, as shown in Figure 46. In addition, you can specify up to 62 characters of information in the entry. The directory entries are arranged in alphameric collating sequence by name.

The track address of each member is recorded by the system as a relative track address within the data set rather than as an absolute track address. Thus, an entire data set can be moved without changing the relative track addresses in the directory. The data set can be considered as one continuous set of tracks regardless of how the space was actually allocated.

If there is not sufficient space available in the directory for an additional entry, or not enough space available within the data set for an additional member, no new members can be stored.

## Partitioned Data Set Directory

The directory of a partitioned data set occupies the beginning of the area allocated to the data set on a direct-access volume. It is searched and maintained by the FIND and STOW macro instructions. The directory consists of member entries arranged in ascending order according to the binary value of the member name or alias.

Member entries vary in length and are blocked into 256-byte blocks. Each block contains as many complete entries as will fit in a maximum of 254 bytes; any remaining bytes are left unused and are ignored. Each directory block contains a 2-byte count field that specifies the number of active bytes in a block (including the count field). As shown in Figure 47, each block is preceded by a hardware-defined key field containing the name of the last member entry in the block, that is, the member name with the highest binary value.

Each member entry contains a member name or alias. There can be as many as 16 aliases for each member. Each entry also contains the relative track address of the member and a count field, as shown in Figure 48. In addition, it may contain a user data field. The last entry in the last directory block has a name field of maximum binary value—all 1s.

NAME
  specifies the member name or alias. It contains up to 8 alphameric characters, left-justified and padded with blanks if necessary.

TTR
  is a pointer to the first block of the member; TT is the number of the track, relative to the beginning of the data set, and R is the number of the block, relative to the beginning of that track.

  Note: This pointer is created by adding 1 to the TTR for the last block of the previous member (which is an end-of-file mark). If track TT is full, the next block will begin at record 1 of track TT + 1, and the pointer will be updated accordingly. The control program finds the block by searching in multitrack mode using TT(R-1) as a



Figure 47. A Partitioned Data Set Directory Block

Member Name | TTR | C | Optional User Data: TTRN | TTRN | TTRN

8 | 3 | | 0-31 halfwords (Maximum 62 bytes)

Pointer to First Record of Member

| 1 If Name is an Alias | Number of User Data TTRNs | Number of User Data Halfwords |

Bits 0 | 1-2 | 3-7
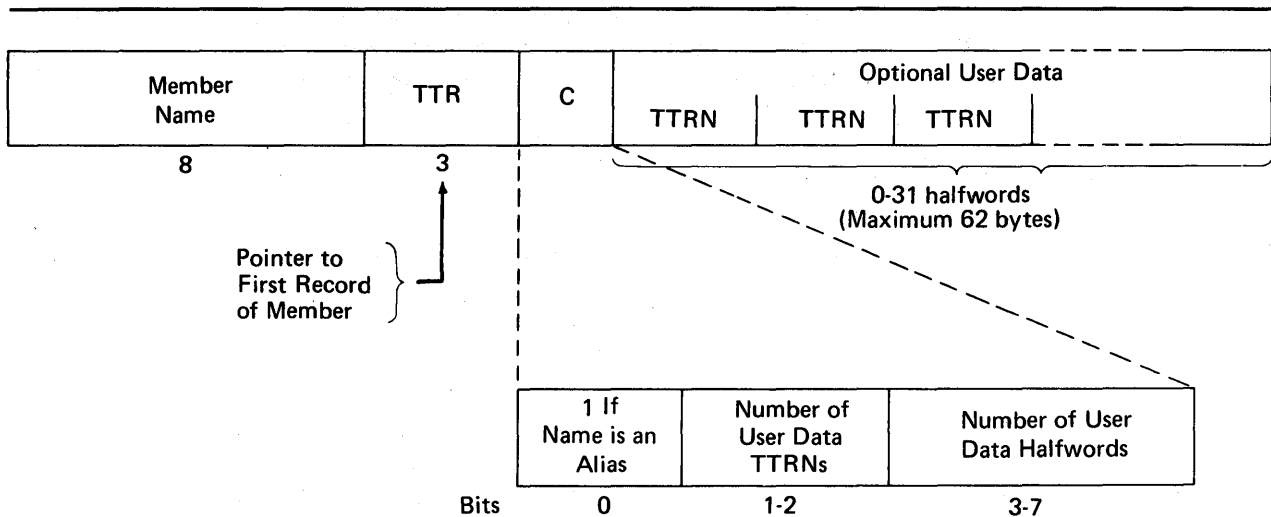
Figure 48. A Partitioned Data Set Directory Entry

search argument.

**C**

specifies the number of halfwords contained in the user data field. It may also contain additional information about the user data field, as shown below:

Bits 0 1-2 3-7

0      when set to 1, indicates that the NAME field contains an alias.

1-2   specifies the number of pointers to locations within the member.

A maximum of three pointers is allowed in the user data field. Additional pointers may be contained in a record referred to as a note list, discussed below. The pointers can be updated automatically if the data set is moved or copied by a utility program such as IEHMOVE. The data set must be marked unmovable under the following conditions:

- More than three pointers are used in the user data field.

- The pointers in the user data field or note list do not conform to the standard format.

- The pointers are not placed first in the user data field.

- Any direct access address (absolute or relative) is embedded in any data blocks or in another data set that refers to this data set.

3-7   contains a binary value indicating the number of halfwords of user data. This number must include the space used by pointers in the user data field.

You can use the user data field to provide variable data as input to the STOW macro instruction. If pointers to locations within the member are provided, they must be 4 bytes long and placed first in the user data field. The user data field format is as follows:

User Data

| TTRN | TTRN | TTRN | Optional |
|------|------|------|----------|

TT    is the relative track address of the note list or area to which you are pointing.

R    is the relative block number on that track.

N    is a binary value that indicates the number of additional pointers contained in a note list pointed to by the TTR. If the pointer is not to a note list, N=0.

A note list consists of additional pointers to blocks within the same member of a partitioned data set. You can divide a member into subgroups and store a pointer to the beginning of each subgroup in the note list. The member may be a load module containing many control sections (CSECTs), each CSECT being a subgroup pointed to by an entry in the note list. You get the pointer to the beginning of the subgroup by using the NOTE macro instruction after you write the first record of the subgroup. Remember that the pointer to the first record of the member is stored in the directory entry by the system.

If the existence of a note list was indicated as shown above, the list can be updated automatically when the data set is moved or copied by a utility program such as IEHMOVE. Each 4-byte entry in the note list has the following format:

| TTRX |
|------|

TT    is the relative track address of the area to which you are pointing.

R    is the relative block number on that track.

X    is available for any use.

To place the note list in the partitioned data set, you must use the WRITE macro instruction. After checking the write operation, use the NOTE macro instruction to determine the address of the list and place that address in the user data field of the directory entry.

## Processing a Member of a Partitioned Data Set

Because a member of a partitioned data set is sequentially organized, it is processed in the same manner as a sequential data set. Either the basic or queued access technique can be used. However, you cannot alter the directory when using the queued technique.

To locate a member or to process the directory, several macro instructions are provided by the operating system. The BLDL macro instruction can be used to structure a list of directory entries in virtual storage; the FIND macro instruction locates a member of the data set for subsequent processing; the STOW macro instruction adds, deletes, replaces, or changes a member name in the directory. To use these macro instructions, you must specify DSORG=PO or POU in the DCB macro instruction. Before issuing a FIND, BLDL, or STOW macro instruction, you must check all preceding input/output operations for completion.

## BLDL—Construct a Directory Entry List

The BLDL macro instruction is used to place directory information in virtual storage. The data is placed in a build list, which you construct before the BLDL macro instruction is issued. The format of the list is similar to that of the directory. For each member name in the list, the system supplies the address of the member and any additional information contained in the directory entry. Note that if there is more than one member name in the list, the member names must be in collating sequence regardless of whether the members are from the same library or from different libraries.

You can optimize retrieval time by directing a subsequent FIND macro instruction to the build list rather than the directory to locate the member to be processed.

The build list, as shown in Figure 49, must be preceded by a 4-byte list description that indicates the number of entries in the list and the length of each entry (12 to 76 bytes). The first 8 bytes of each entry contain the member name or alias. The next 6 bytes must be available to contain the starting address of the member plus some control data. If there is no user data entry, only the TTR and C fields are required. If additional information is to be supplied from the directory, up to 62 bytes can be reserved.

## FIND—Position to a Member

To determine the starting address of a specific member, you must issue a FIND macro instruction. The system places the correct address in the data control block so that a subsequent input or output operation begins processing at that point.

There are two ways you can direct the system to the right member when you use the FIND macro instruction: specify the address of an area containing the name of the member or specify the address of the TTR field of the entry in a build list you have created by using the BLDL macro instruction. In the first case, the system searches the directory of the data set for the relative track address; in the second case, no search is required because the relative track address is in the build list entry.



Programmer Supplies:

| | |
|---|---|
| FF | Number of member entries in list. |
| LL | Even number giving byte length of each entry (minimum of 12). |
| Member name | Eight bytes, left-justified. |

BLDL Supplies:

| | |
|---|---|
| TTR | Member starting location. |
| K | If only data set = 0. If concatenation = number. Not required if no user data. |
| Z | Source of directory entry. Private library = 0. Link library = 1. Job or step library = 2. Not required if no user data. |
| C | Same C field from directory. Gives number of user data halfwords. |
| User data | As much as will fit in entry. |

Figure 49. Build List Format

If you want to process only one member, you can process it as a sequential data set (DSORG=PS) using either BSAM or QSAM. You indicate the name of the member you want to process and the name of the partitioned data set in the DSNAME parameter of the DD statement. When you open the data set, the system places the starting address in the data control block so that a subsequent GET or READ macro instruction begins processing at that point. You cannot use the FIND, BLDL, or STOW macro instructions when you are processing one member as a sequential data set.

### STOW—Update the Directory

When you add several members to a partitioned data set, you must issue a STOW macro instruction after writing each member so that an entry for each one will be added to the directory. To use the STOW macro instruction, DSORG=PO or POU must be specified in the DCB macro instruction.

You can also use the STOW macro instruction to delete, replace, or change a member name in the directory, as well as to store additional information with the directory entry. Since an alias can also be stored in the directory the same way, you should be consistent in altering all names associated with a given member. For example, if you replace a member, you must delete related alias entries or change them so that they point to the new member.

If you add only one member to a partitioned data set and indicate the member name in the DSNAME parameter of the DD statement, it is not necessary for you to use BPAM and a STOW macro instruction in your program. If you wish to do so, you may use BPAM and STOW, or BSAM or QSAM. If you use a sequential access method, or if you use BPAM and issue a CLOSE macro instruction without issuing a STOW macro instruction, the system will issue a STOW macro instruction using the member name you have specified on the DD statement. When the system issues the STOW, the directory entry that is added is the minimum length (12 bytes). This automatic STOW macro instruction will not be issued if the CLOSE macro instruction is a TYPE=T. The DISP parameter on the DD statement determines what *directory action* parameter will be chosen by the system for the STOW macro instruction.

If DISP=NEW or MOD was specified, a STOW macro instruction with the add option will be issued. If the member name on the DD statement is not present in the data set directory, it will be added. If the member name is already present in the directory, the task will be abnormally terminated.

If DISP=OLD was specified, a STOW macro instruction with the replace option will be issued. The member name will be inserted into the directory, either as an addition if the name is not already present or as a replacement if the name is present.

Thus, with an existing data set, you should use DISP=OLD to force a member into the data set; you should use DISP=MOD to add members with protection against the accidental destruction of an existing member.

## Creating a Partitioned Data Set

If you have no need to add entries to the directory, that is, the STOW and BLDL macro instructions will not be used, you can create a new data set and write the first member as follows (see Figure 50):

- Code DSORG=PS or DSORG=PSU in the DCB macro instruction.

- Indicate in the DD statement that the data is to be stored as a member of a new partitioned data set, that is, DSNAME=name (membername) and DISP=NEW.

- Request space for the member and the directory in the DD statement.

- Process the member with an OPEN macro instruction, a series of PUT or WRITE macro instructions, and then a CLOSE macro instruction. A STOW macro instruction is issued automatically when the data set is closed.

As a result of these steps, the data set and its directory are created, the records of the member are written, and a 12-byte entry is made in the directory.

To add additional members to the data set, follow the same procedure. However, a separate DD statement (with the space request omitted) is required for each member. The disposition should be specified as modify, DISP=MOD. The data set must be closed and reopened each time a new member is specified.

To take full advantage of the STOW macro instruction, and thus the BLDL and FIND macro instructions in future processing, you can provide additional information with each directory entry. You do this by using the basic access technique, which also allows you to process more than one member without closing and reopening the data set, as follows (see Figure 51):

- Request space in the DD statement for the members and the directory.

- Define DSORG=PO or DSORG=POU in the DCB macro instruction.

- Use WRITE and CHECK to write and check the member records.

- Use NOTE to note the location of any note list written within the member, if there is a note list.

---

```
//PDSDD      DD           ---,DSNAME=MASTFILE(MEMBERK),SPACE=(TRK,(100,5,7)),        C
                          DISP=(NEW,KEEP)
           ...
OUTDCB      DCB     --,DSORG=PS,DDNAME=PDSDD,---
           ...
            OPEN         (OUTDCB,(OUTPUT))
            PUT[or WRITE]
           ..
            CLOSE        (OUTDCB)                          Automatic Stow
           ...
```

Figure 50. Creating One Member of a Partitioned Data Set

---

```
//PDSDD      DD           --,DSNAME=MASTFILE,SPACE=(TRK,(100,5,7)),DISP=MOD
             ...
OUTDCB       DCB          --,DSORG=PO,DDNAME=PDSDD,--
             OPEN         (OUTDCB,(OUTPUT))
             WRITE        **          Write and check first record of member.
             CHECK                    The system will supply the relative
*                                     track address for the directory entry.
             WRITE                    Write and check remaining records of
             CHECK                    member.
*
             NOTE                     If you are dividing the member into
             ST                       subgroups, note the location of the first
*                                     record in subgroup, storing pointer
*                                     in note list.
             WRITE                    Write note list at end of member.
             CHECK
             NOTE                     Note location of note list, storing
             ST                       pointer in list for STOW.
             STOW                     Enter information in directory for
*                                     this member after all records and note
*                                     lists are written.
Repeat from ** for each additional member
             CLOSE (OUTDCB)
```

Figure 51. Creating Members of a Partitioned Data Set Using STOW

- When all the member records have been written, issue a STOW macro instruction to enter the member name, its location pointer, and any additional data in the directory.

- Continue to write, check, note, and stow until all the members of the data set and the directory entries have been written.

## Retrieving a Member of a Partitioned Data Set

To retrieve a specific member from a partitioned data set, either the basic or queued access technique can be used as follows (see Figure 52):

- Code DSORG=PS or DSORG=PSU in the DCB macro instruction.

- Indicate in the DD statement that the data is a member of an existing partitioned data set by coding DSNAME=name(membername) and DISP=OLD.

- Process the member with an OPEN macro instruction, a series of GET and READ macro instructions, and then a CLOSE macro instruction.

```
//PDSDD      DD           --,DSNAME=MASTFILE(MEMBERK),DISP=OLD
             ...
INDCB        DCB          --,DSORG=PS,DDNAME=PDSDD,--
             OPEN   (INDCB)                            Automatic Find
             GET (or READ)
             CLOSE (INDCB)
             ...
```

Figure 52. Retrieving One Member of a Partitioned Data Set

When your program is executed, the directory is searched automatically and the location of the member is placed in the DCB.

To process several members without closing and reopening, or to take advantage of additional data in the directory, this technique should be used (see Figure 53):

- Code DSORG=PO or POU in the DCB macro instruction.

- Build a list (BLDL) of needed member entries from the directory.

- Indicate in the DD statement the data set name of the partitioned data set by coding DSNAME=name and DISP=OLD.

- Use the FIND or POINT macro instruction to prepare for reading the member records.

- The records may be read from the beginning of the member, or a note list may be read first, to obtain additional locations that point to subcategories within the member.

- Read (and check) the records until all those required have been processed.

- Point to additional categories, if required, and read the records.

- Repeat this procedure for each member to be retrieved.

```
//PDSDD      DD            --,DSNAME=MASTFILE,DISP=OLD
            ...
INDCB       DCB           --,DSORG=PO,DDNAME=PDSDD,--
            OPEN          (INDCB)
            BLDL                      Build a list of selected member names
                                      in virtual storage.
            FIND (or POINT)
**Read note list.
            READ
            CHECK
            POINT                     Locate subgroup by using note list.
            READ
            CHECK                     Read member records.
```

Repeat from ** for each additional member.

```
            CLOSE         (INDCB)
            ...
```

Figure 53. Retrieving Several Members of a Partitioned Data Set Using BLDL, FIND, and POINT

## Updating a Member of a Partitioned Data Set

A member of a partitioned data set can be updated in place, or can be deleted and rewritten as a new member.

### Updating in Place

When you update in place, you read records, process them, and write them back to their original positions without destroying the remaining records on the track. The following rules apply:

- You must specify the update option (UPDAT) in the OPEN macro instruction. To perform the update, you can use only the READ, WRITE, CHECK, NOTE, POINT, FIND, and BLDL macro instructions.

- You cannot update concatenated partitioned data sets.

- You cannot use chained scheduling.

- You cannot delete any record or change its length; you cannot add new records.

A record must be retrieved by a READ macro instruction before it can be updated by a WRITE macro instruction. Both macro instructions must be execute forms that refer to the same DECB; the DECB must be provided by a list form. (The execute and list forms of the READ and WRITE macro instructions are described in *OS/VS1 Data Management Macro Instructions.*)

**Updating With QSAM:** You can update a member of a partitioned data set using the locate mode of QSAM (DCB specifies MACRF=PL) and using the PUTX macro instruction. The DD statement must specify the data set and member name in the DSNAME parameter. This method allows only the updating of the member specified in the DD statement.

**Updating With Overlapped Operations:** To overlap input/output and CPU activity, you can start several read or write operations before checking the first for completion. You cannot overlap read and write operations, however, as operations of one type must be checked for completion before operations of the other type are started or resumed. Note that each concurrent read or write operation requires a separate channel program and a separate DECB. If a single DECB were used for successive read operations, only the last record read could be updated.

In Figure 54, overlap is achieved by having a read or write request outstanding while each record is being processed. Note the use of the execute and list forms of the READ and WRITE macro instructions, identified by the operands MF=E and MF=L.

## Rewriting a Member

There is no actual update option that can be used to add or extend records in a partitioned data set. If you want to extend or add a record within a member, you must rewrite the complete member in another area of the data set. Since space is allocated when the data set is created, there is no need to request additional space. Note, however, that a partitioned data set must be contained on one volume. If sufficient space has not

```
//PDSDD   DD        DSNAME=MASTFILE(MEMBERK),DISP=OLD,---

         ...
UPDATDCB DCB       DSORG=PS,DDNAME=PDSDD,MACRF=(R,W),NCP=2,EODAD=FINISH
         READ      DECBA,SF,UPDATDCB,AREAA,MF=L          Define DECBA
         READ      DECBB,SF,UPDATDCB,AREAB,MF=L          Define DECBB
AREAA    DS        ---                                   Define buffers
AREAB    DS        ---
         ...
         OPEN      (UPDATDCB,UPDAT)                      Open for update
         LA        2,DECBA                               Load DECB addresses
         LA        3,DECBB
READRECD READ      (2),SF,MF=E                           Read a record
NEXTRECD READ      (3),SF,MF=E                           Read the next record
         CHECK     (2)                                   Check previous read operation
         (If update is required, branch to R2UPDATE)

         LR        4,3                                   If no update is required,
         LR        3,2                                   switch DECB addresses in
         LR        2,4                                   registers 2 and 3
         B         NEXTRECD                              and loop
```

*In the following statements, "R2" and "R3" refer to the records that were read using the DECBs whose addresses are in registers 2 and 3, respectively. Either register may point to either DECBA or DECBB.*

```
R2UPDATE CALL      UPDATE,((2))                          Call routine to update R2
         CHECK     (3)                                   Check read for next record (R3)
         WRITE     (2),SF,MF=E                           Write updated R2
         (If R3 requires an update, branch to R3UPDATE)
         CHECK     (2)                                   If R3 requires no update, check
         B         READRECD                              write for R2 and loop
R3UPDATE CALL      UPDATE,((3))                          Call routine to update R3
         WRITE     (3),SF,MF=E                           Write updated R3
         CHECK     (2)                                   Check write for R2
         CHECK     (3)                                   Check write for R3
         B         READRECD                              Loop
FINISH   CLOSE     (UPDATDCB)                            End-of-Data exit routine
         ...
```

Figure 54. Updating a Member of a Partitioned Data Set

been allocated, the data set must be reorganized by the IEBCOPY utility program.

When you rewrite the member, you must provide two DCBs, one for input and one for output. Both DCB macro instructions can refer to the same data set, that is, only one DD statement is required.

You can reflect the change in location of the member either automatically, by indicating a disposition of OLD, or by using the STOW macro instruction. Although the old member is, in effect, deleted, its space cannot be reused until the data set is reorganized.

# Processing an Indexed Sequential Data Set

The organization of an indexed sequential data set allows you a great deal of flexibility in the operations you can perform. The data set can be read or written sequentially, individual records can be processed in any order, records can be deleted, and new records can be added. The system automatically locates the proper position in the data set for new records and makes any necessary adjustments when records are deleted. However, when accessing an indexed sequential data set, the key of user control blocks and user storage must be in the same protection key as that specified in the TCB (TCBPKF).

The queued access technique must be used to create an indexed sequential data set. It can also be used to sequentially process or update the data set and to add records to the end of the data set. The basic access technique can be used to insert new records between records already in the data set and to update the data set directly.

## *Indexed Sequential Data Set Organization*

The records in an indexed sequential data set are arranged according to collating sequence by a *key field* in each record. Each block of records is preceded by a key field that corresponds to the key of the last record in the block.

An indexed sequential data set resides on direct-access storage devices and can occupy up to three different areas:

- *Prime Area*—This area, also called the prime data area, contains data records and related track indexes. It exists for all indexed sequential data sets.

- *Overflow Area*—This area contains records that overflow from the prime area when new data records are added. It is optional.

- *Index Area*—This area contains master and cylinder indexes associated with the data set. It exists for a data set that has a prime area occupying more than one cylinder.

The indexes of an indexed sequential data set are analogous to the card catalog in a library. For example, if the library user knows the name of the book or the author, he can look in the card catalog and obtain a catalog number that will enable him to locate the book in the book files. He would then go to the shelves and proceed through rows until he found the shelf containing the book. Usually each row contains a sign to indicate the beginning and ending numbers of all books in that particular row. Thus, as he proceeded through the rows, he would compare the catalog number obtained from the index with the numbers posted on each row. Upon locating the proper row, he would then search that row for the shelf that contained the book. Then he would look at the individual book numbers on that shelf until he found the particular book.

ISAM uses the indexes in much the same way to locate records in an indexed sequential data set.

As the records are written in the prime area of the data set, the system accounts for the records contained on each track in a *track index area*. Each entry in the track index identifies the key of the last record on each track. There is a track index for each cylinder in the data set. If more than one cylinder is used, the system develops a higher-level index called a *cylinder index*. Each entry in the cylinder index identifies the key of the last record in the cylinder. To increase the speed of searching the cylinder index, you can request that a *master index* be developed for a specified number of cylinders, as shown in Figure 55.

Rather than reorganize the whole data set when records are added, you can request that space be allocated for additional records in an overflow area.

## Prime Area

Records are written in the prime area when the data set is created or updated. The last track of prime data is reserved for an end-of-file mark. The portion of Figure 55 labeled Cylinder 1 illustrates the initial structure of the prime area. Although the prime area can extend across several noncontiguous areas of the volume, all the records are written in key sequence. Each record must contain a key; the system automatically writes the key of the highest record before each block.

When the ABSTR option of the SPACE parameter of the DD statement is used to generate a multivolume prime area, the VTOC of the second volume and on all succeeding volumes must be contained within cylinder 0 of the volume.

## Index Areas

The operating system generates track and cylinder indexes automatically. Up to three levels of master indexes are created if requested.
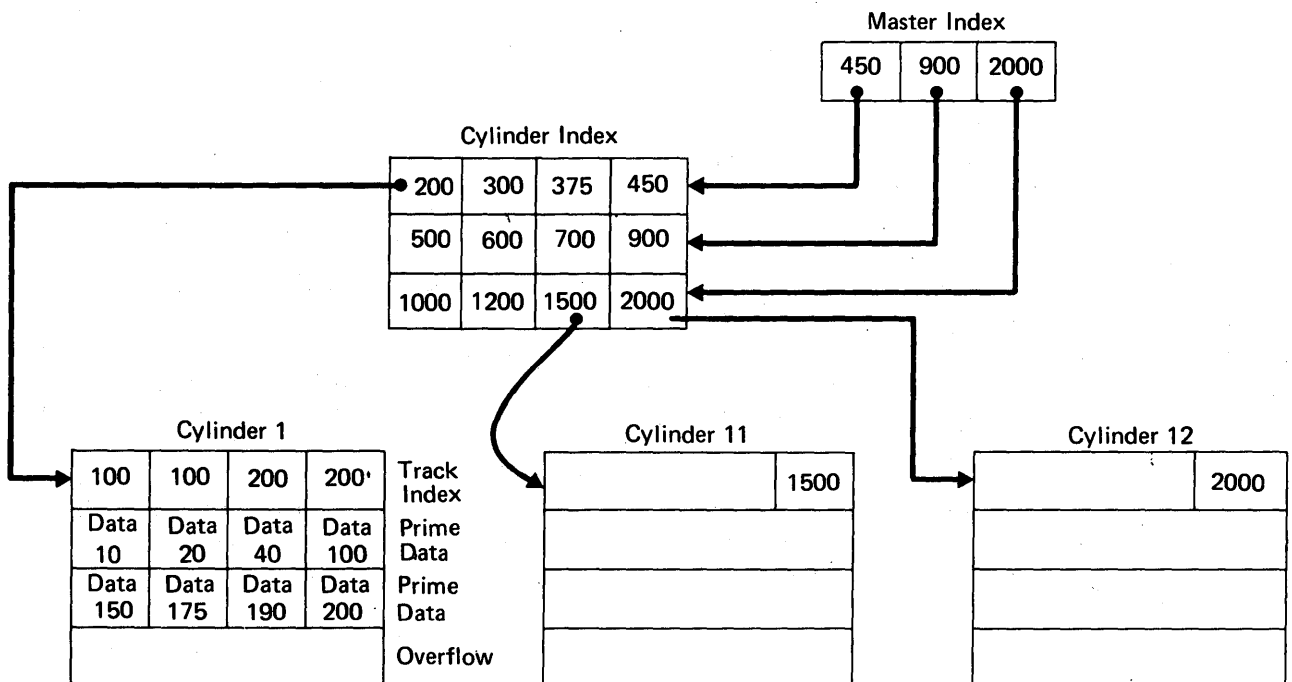


Figure 55. Indexed Sequential Data Set Organization

**Track Index:** This is the lowest level of index and is always present. There is one track index for each cylinder in the prime area; it is written on the first track(s) of the cylinder that it indexes.

The index consists of a series of paired entries, that is, of a normal entry and an overflow entry for each prime track. For fixed-length records, each normal entry (and also DCBFIRSH) points to either record 0 or the first prime record on a shared track. For variable-length records, the normal entry contains the key of the highest record on the track and the address of the last record on the track. The overflow entry is originally the same as the normal entry. (This is why 100 appears twice on the track index for cylinder 1 in Figure 55.) The overflow entry is changed when records are added to the data set. Then the overflow entry contains the key of the highest overflow record and the address of the lowest overflow record logically associated with the track. Figure 56 shows the format of a track index.

If all the tracks allocated for the prime data area are not used, the index entries for the unused ones are flagged as inactive. The last entry of each track index is a dummy entry indicating the end of the index. When fixed-length record format has been specified, the remainder of the last track of each cylinder used for a track index contains prime data records if there is room for them.

Each index entry has the same format. It is an unblocked, fixed-length record consisting of a count, a key, and a data area. The length of the key corresponds to the length of the key area in the record to which it points. The data area is always 10 bytes long. It contains the full address of the track or record to which the index points, as well as the level of the index and the entry type.

**Cylinder Index:** For every track index created, the system generates a cylinder index entry. There is one cylinder index for a data set, each entry of which points to a track index. Since there is one track index per cylinder, there is one cylinder index entry for each cylinder in the prime data area, except in the case of a 1-cylinder prime area. As with track indexes, inactive entries are created for any unused cylinders in the prime data area.
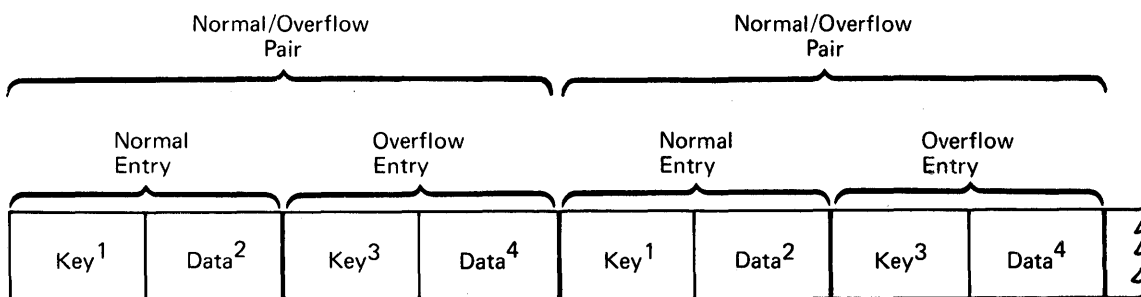


[1]Normal key    = key of the highest record on the prime data track

[2]Normal data   = address of the prime data track

[3]Overflow key  = key of the highest overflow record logically associated with the prime data track

[4]Overflow data = address of the lowest overflow record logically associated with the prime data track

Notes:

- If there are no overflow records, overflow key and data entries are the same as normal key and data entries.
- This figure is a logical representation only; that is, it makes no attempt to show the physical size of track index entries.

Figure 56. Format of Track Index Entries

**Master Index:** As an optional feature, the operating system creates, at your request, a master index. The presence of this index makes long, serial searches through a large, cylinder index unnecessary.

You can specify the conditions under which you want a master index created. For example, if you have specified NTM=3 and OPTCD=M in your DCB macro instruction, a master index is created when the cylinder index exceeds 3 tracks. The master index consists of one entry for each track of cylinder index. If your data set is extremely large, a higher-level master index is created when the first-level master index exceeds three tracks. This higher-level master index consists of one entry for each track of the first-level master index. This procedure can be repeated for as many as three levels of master index.

## Overflow Areas

As records are added to an indexed sequential data set, space is required to contain those records that will not fit on the prime data track on which they belong. You can request that a number of tracks be set aside as a *cylinder overflow area* to contain overflows from prime tracks in each cylinder. An advantage of using cylinder overflow areas is a reduction of search time required to locate overflow records. A disadvantage is that there will be unused space if the additions are unevenly distributed throughout the data set.

Instead of, or in addition to, cylinder overflow areas, you can request an *independent overflow area*. Overflow from anywhere in the prime data area is placed in a specified number of cylinders reserved solely for overflow records. An advantage of having an independent overflow area is a reduction in unused space reserved for overflow. A disadvantage is the increased search time required to locate overflow records in an independent area.

If you request both cylinder overflow and independent overflow, the cylinder overflow area is used first. It is a good practice to request cylinder overflow areas large enough to contain a reasonable number of additional records and an independent overflow area to be used as the cylinder overflow areas are filled.

## *Adding Records to an Indexed Sequential Data Set*

Either the queued access technique or the basic access technique may be used to add records to an indexed sequential data set. A record to be inserted between records already in the data set must be inserted by the basic access method using WRITE KN (key new). Records added to the end of a data set, that is, records with successively higher keys, may be added to the prime data area or the overflow area by the basic access method using WRITE KN, or they may be added to the prime data area by the queued access technique using the PUT macro instruction.

### Inserting New Records into an Existing Indexed Sequential Data Set

As you add records to an indexed sequential data set, the system inserts each record in its proper sequence according to the record key. The remaining records on the track are then moved up one position each. If the last record does not fit on the track, it is written in the first available location in the overflow area. A 10-byte *link field* is added to the record put in the overflow area to connect it logically to the correct track. The proper adjustments are made to the track index entries. This procedure is illustrated in Figure 57.

Subsequent additions are written either on the prime track or as part of the *overflow chain* from that track. If the addition belongs after the last prime record on a track but before a previous overflow record from that track, it is written in the first available location in the overflow area. Its link field contains the address of the next record in the chain.
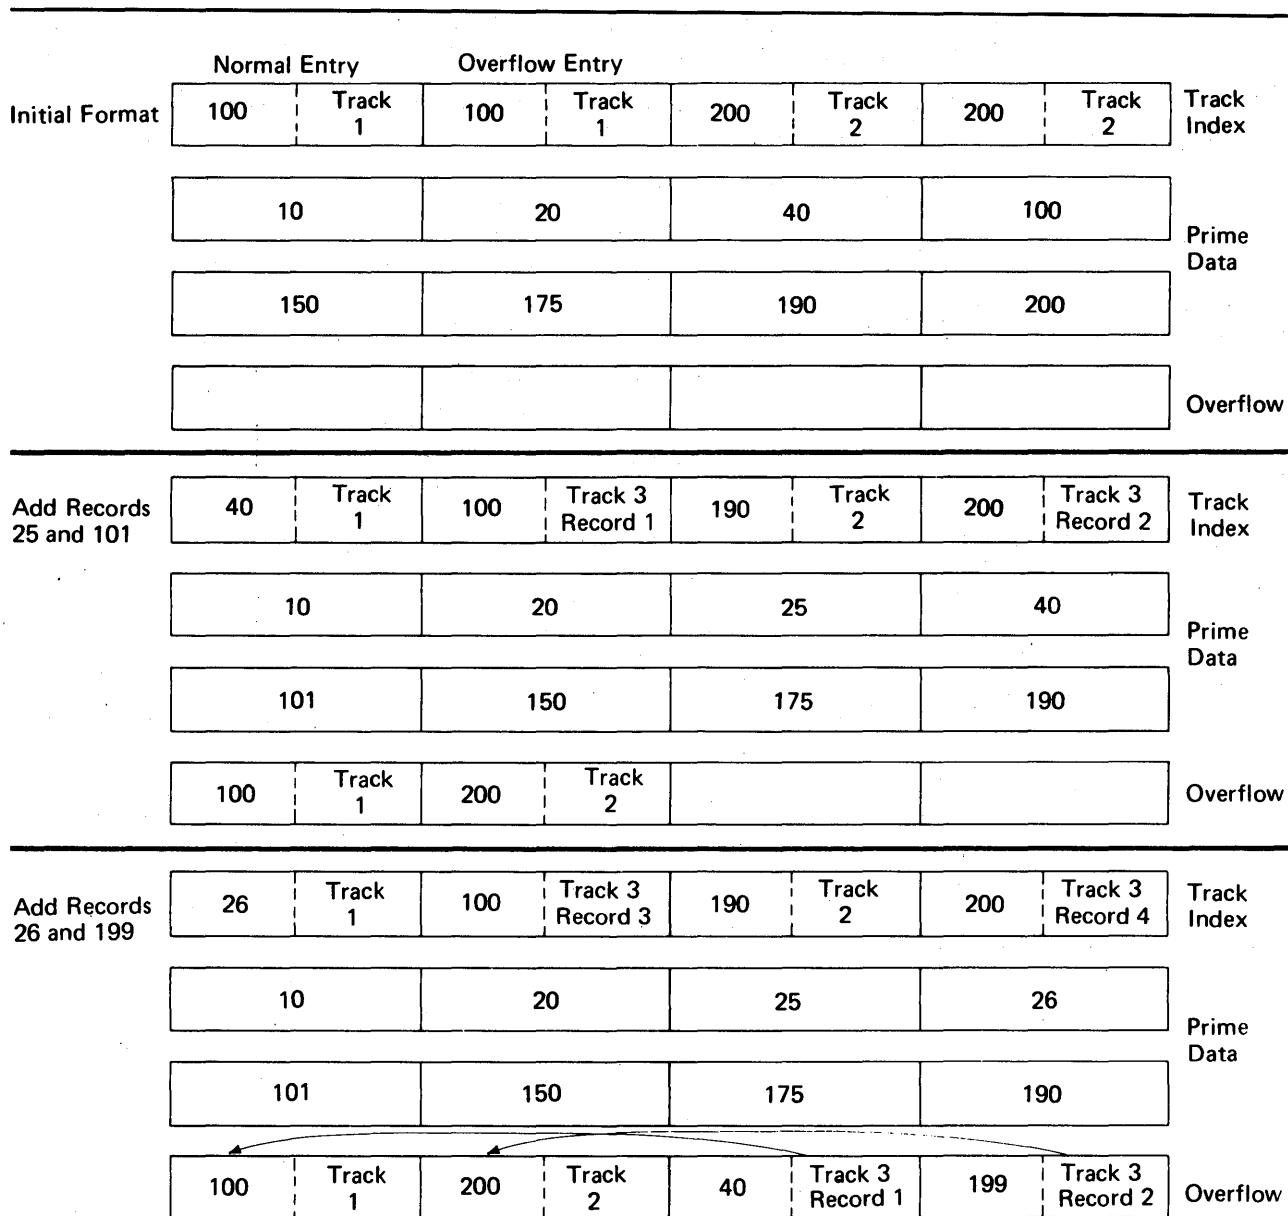
**Figure 57. Adding Records to an Indexed Sequential Data Set**

Normal Entry     Overflow Entry

**Initial Format**

| 100 | Track 1 | 100 | Track 1 | 200 | Track 2 | 200 | Track 2 | Track Index |

| 10 | 20 | 40 | 100 | Prime Data |
| 150 | 175 | 190 | 200 | |

| | | | | Overflow |

**Add Records 25 and 101**

| 40 | Track 1 | 100 | Track 3 Record 1 | 190 | Track 2 | 200 | Track 3 Record 2 | Track Index |

| 10 | 20 | 25 | 40 | Prime Data |
| 101 | 150 | 175 | 190 | |

| 100 | Track 1 | 200 | Track 2 | | | Overflow |

**Add Records 26 and 199**

| 26 | Track 1 | 100 | Track 3 Record 3 | 190 | Track 2 | 200 | Track 3 Record 4 | Track Index |

| 10 | 20 | 25 | 26 | Prime Data |
| 101 | 150 | 175 | 190 | |

| 100 | Track 1 | 200 | Track 2 | 40 | Track 3 Record 1 | 199 | Track 3 Record 2 | Overflow |

## Adding New Records to the End of an Indexed Sequential Data Set

Records added to the end of a data set, that is, records with successively higher keys, may be added by the basic access method using WRITE KN (key new), or by the queued access method using the PUT macro instruction (resume load). In either case records may be added to the prime data area.

When you use the WRITE KN macro instruction, the record being added is placed in the prime data area only if there is room for it on the prime data track containing the record with the highest key currently in the data set. If there is not sufficient room on that track, the record is placed in the overflow area and linked to that prime track even though additional prime data tracks originally allocated have not been filled.

When you use the PUT macro instruction (resume load), records are added to the prime data area until the space originally allocated is filled. Once this allocated prime area is filled, you can add records to the data set using WRITE KN, in which case they will be

placed in the overflow area. Resume load is discussed in more detail later under "Creating an Indexed Sequential Data Set."

In order to add records with successively higher keys using the PUT macro instruction (resume load):

- The key of any record to be added must be higher than the highest key currently in the data set.

- The DD statement must specify DISP=MOD, or the EXTEND option must be specified in the OPEN macro instruction.

- The data set must have been successfully closed when it was created or when records were previously added using the PUT macro instruction.

You may continue to add fixed-length records in this manner until the original space allocated for prime data is exhausted.

When you add records to an indexed sequential data set using the PUT macro instruction (resume load), new entries are also made in the indexes. During resume load on a data set with a partially filled track and/or a partially filled cylinder, the track index entry and/or the cylinder index entry is overlaid when the track or cylinder is filled. If resume load abnormally terminates after these index entries have been overlaid, a subsequent resume load will get a sequence check when adding a key that is higher than the highest key at the last successful CLOSE but lower than the key in the overlaid index entry. When the SYNAD exit is taken for a sequence check, register 0 contains the address of the highest key of the data set.

## Maintaining an Indexed Sequential Data Set

An indexed sequential data set must be reorganized occasionally for two reasons:

- The overflow area will eventually be filled.

- Additions increase the time required to locate records directly.

The frequency of reorganization depends on the activity of the data set and on your timing and storage requirements. There are two ways you can accomplish reorganization:

- You can reorganize the data set in two passes by writing it sequentially into another area of direct-access storage or magnetic tape and then recreating it in the original area.

- You can reorganize the data set in one pass by writing it directly into another area of direct-access storage. In this case, the area occupied by the original data set cannot be used by the reorganized data set.

The operating system maintains statistics that are pertinent to reorganization. The statistics, written on the direct-access volume and available in the DCB for checking, include the number of cylinder overflow areas, the number of unused tracks in the independent overflow area, and the number of references to overflow records other than the first. They appear in the RORG1, RORG2, and RORG3 fields of the DCB.

If you indicate when creating or updating the data set that you want to be able to flag records for deletion during updating, you can set the delete code (the first byte of a fixed-length record or the fifth byte of a variable-length record) to X'FF'. If a flagged record is forced off its prime track during a subsequent update, it will not be rewritten in the overflow area, as shown in Figure 58, unless it has the highest key on that cylinder. Similarly, when you process sequentially, flagged records are not retrieved for processing. During direct processing, flagged records are retrieved like any other records, and you should check them for the delete code.

Note that to use the delete option, RKP must be greater than 0 for fixed-length records and greater than 4 for variable-length records.
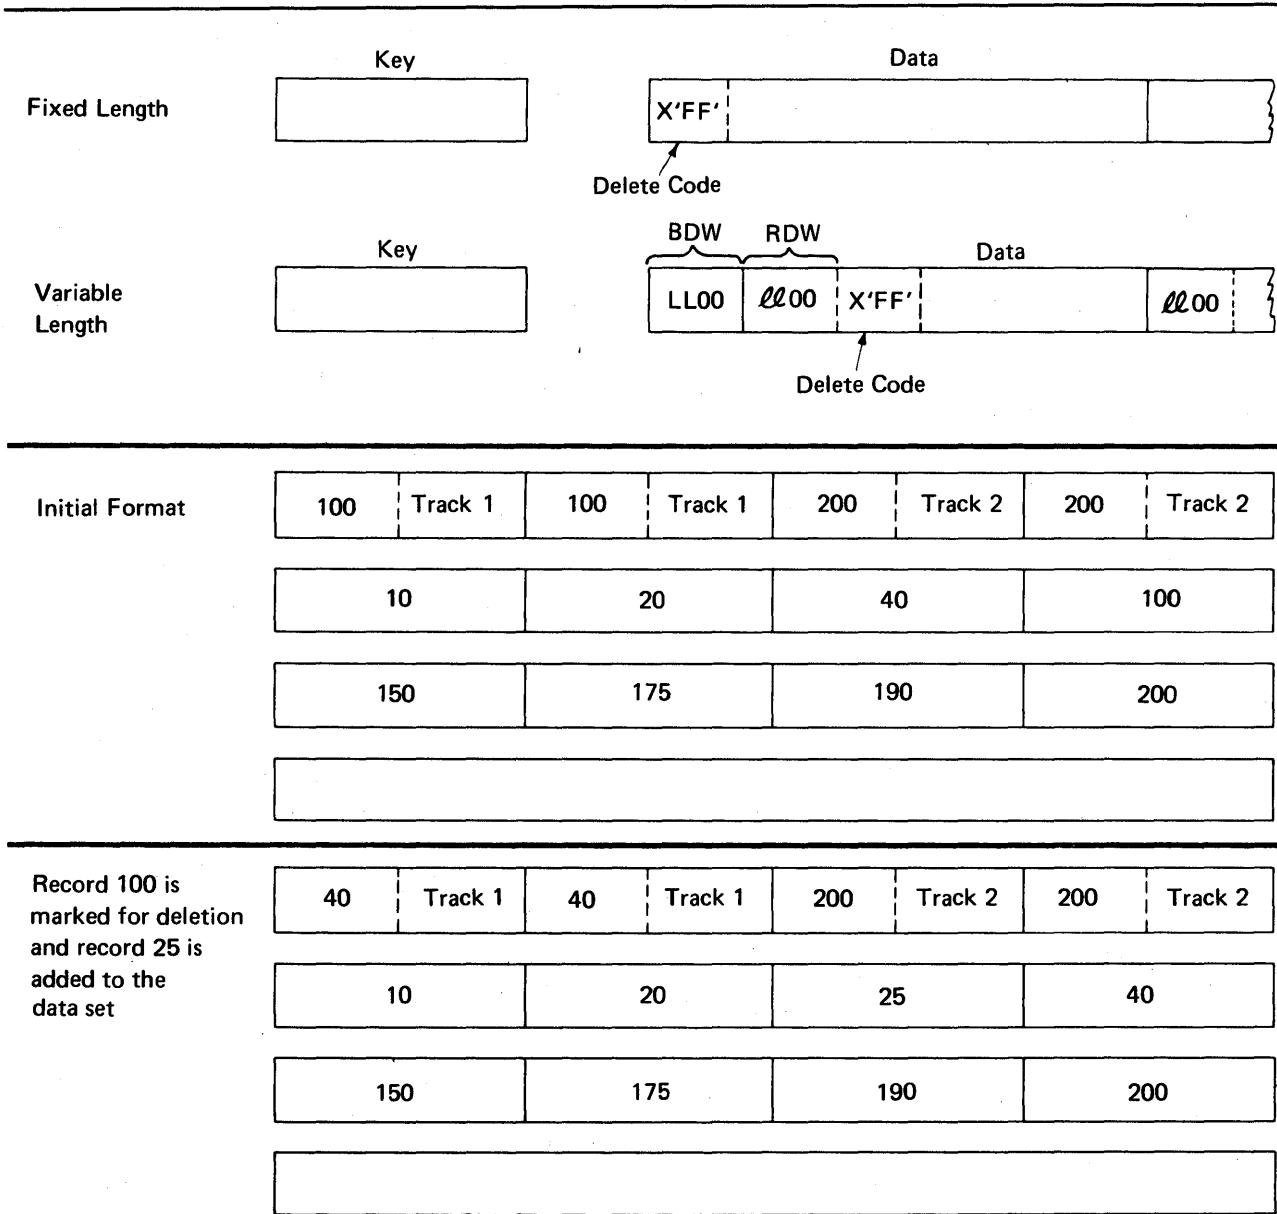


Figure 58. Deleting Records From an Indexed Sequential Data Set

## Indexed Sequential Buffer and Work Area Requirements

The only case in which you will ever have to compute the buffer length (BUFL) requirements for your program is when you use the BUILD or GETPOOL macro instruction to construct the buffer area. If you are creating an indexed sequential data set (using the PUT macro instruction), each buffer must be 8 bytes longer than the blocksize to allow for the hardware count field, that is:

Buffer length = 8 + Blocksize

| (8) | Data (BLKSIZE) |
|-----|----------------|

◄─────────────────────── Buffer ───────────────────────

One exception to this formula arises when you are dealing with an unblocked format-F record whose key field precedes the data field; its relative key position is 0 (RKP=0). In that case the key length must also be added, that is:

Buffer length = 8 + Key length + Record length

| (8) | Key (KEYLEN) | Data (LRECL) |
|-----|--------------|--------------|

◄─────────────────────── Buffer ───────────────────────►

The buffer requirements for using the queued access technique to read or update (using the GET or PUTX macro instruction) an indexed sequential data set are discussed below.

For fixed-length unblocked records when both the key and data are to be read and for variable-length unblocked records, padding is added so that the data will be on a doubleword boundary, that is:

Buffer length = Key length + Padding + 10 + Blocksize

| Key (KEYLEN) | Padding | Link (10) | Data (BLKSIZE) |
|--------------|---------|-----------|----------------|

◄─────────────────────── Buffer ───────────────────────►

For fixed-length unblocked records when only data is to be read:

Buffer length = 16 + LRECL

| Padding (6) | Link (10) | Data (LRECL) |
|-------------|-----------|--------------|

◄─────────────────────── Buffer ───────────────────────►

For fixed-length blocked records:

Buffer length = 16 + Blocksize

| Padding (6) | Link (10) | Data (BLKSIZE) |
|---|---|---|

◄─────────────────── Buffer ───────────────────►

For variable-length blocked records, padding is 2 if the buffer starts on a fullword boundary that is not also a doubleword boundary or 6 if the buffer starts on a doubleword boundary, that is:

Buffer length = 12 or 16 + Blocksize

| Padding | Link (10) | Data (BLKSIZE) |
|---|---|---|

◄─────────────────── Buffer ───────────────────►

If you are using the input data set with fixed-length, unblocked records as a basis for creating a new data set, a work area is required.

The size of the work area is given by:

Work area = Key length + Record length

| Key | Data (LRECL) |
|---|---|

◄─────────────────── Work Area ───────────────────►

If you are reading only the data portion of fixed-length unblocked records or variable-length records, the work area is the same size as the record, that is:

Work area = Record length

| Data (LRECL) |
|---|

◄─────────────────── Work Area ───────────────────►

When you use the basic access technique to update records in an indexed sequential data set, the key length field need not be considered in determining your buffer requirements. The area for fixed-length records must be:

Buffer length = 16 + Blocksize

| Padding (6) | Link (10) | Data (BLKSIZE) |
|---|---|---|

◄─────────────────── Buffer ───────────────────►

For variable-length records, padding is 2 if the buffer starts on a fullword boundary that is not also a doubleword boundary or 6 if a buffer starts on a doubleword boundary. Thus, the area must be:

Buffer length = 12 or 16 + Blocksize

| Padding | Link (10) | Data (BLKSIZE) |
|---------|-----------|----------------|

◄────────────────────── Buffer ──────────────────────►

You can speed up the process of adding fixed-length or variable-length records to a data set by using the MSWA parameter of the DCB macro instruction to provide a special work area for the operating system. The size of the work area (SMSW parameter in the DCB) must be large enough to contain a full track of data, the count fields of each block, and the work space for inserting the new record.

The size of the work area needed varies according to the record format and the device type. You can calculate it during execution using device-dependent information obtained with the DEVTYPE macro instruction and data set information from the DSCB obtained with the OBTAIN macro instruction. The DEVTYPE and OBTAIN macro instructions are discussed in *OS/VS1 Data Management for System Programmers*.

Note that you can use the DEVTYPE macro instruction only if the index and prime areas are on devices of the same type or if the index area is on a device with a larger track capacity than that of the device containing the prime area. If you are not trying to maintain device independence, you may precalculate the size of the work area needed and specify it in the SMSW field of the DCB macro instruction. The maximum value for SMSW is 65,535.

For calculating the size of the work area, refer to the storage device capacities shown in Figure 66 under "Estimating Space Requirements" and the device overhead formulas given in the same section.

For fixed-length blocked records, SMSW is calculated as follows:

SMSW = HIRPD(BLKSIZE + 8) + LRECL + KEYLEN

The formula for fixed-length unblocked records is

SMSW = HIRPD(KEYLEN + LRECL + 8) + 2

The value for HIRPD is in the index (format-2) DSCB. *OS/VS1 System Data Areas* shows the exact location of this field in the index DSCB. If you don't use the MSWA and SMSW parameters, the control program supplies a work area using the formula BLKSIZE + LRECL + KEYLEN.

For variable-length records, SMSW may be calculated by one of two methods. The first method may lead to faster processing although it may require more storage than the second method.

The first method is as follows:

SMSW = HIRPD(BLKSIZE + 8) + LRECL + KEYLEN + 10

The second method is as follows:

$$SMSW = \frac{(Track\ Capacity - Bn + 1)}{Bi}(BLKSIZE) + 8(HIRPD) + LRECL + KEYLEN + 10 + (REM - N - KEYLEN)$$

In all of the above formulas, the terms BLKSIZE, LRECL, KEYLEN, and SMSW are the same as the parameters in the DCB macro instruction. REM is the remainder of the

division operation in the formula and N is the first constant in the Bi formulas described in Figure 67. (REM-N-KEYLEN) is added only if it is positive. The second method yields a minimum value for SMSW. Therefore, the first method is valid only if its application results in a value higher than the value that would be derived from the second method. If neither MSWA nor SMSW is specified, the control program supplies the work area for variable-length records, using the second method to calculate the size.

Another technique to increase the speed of processing is to provide space in virtual storage for the highest-level index. To specify the address of this area, use the MSHI operand of the DCB. When the address of this area is specified, you must also specify its size, which you can do by using the SMSI operand of the DCB. The maximum value for SMSI is 65,535. If you do not use this technique, the index on the volume must be searched.

The size of the storage area (SMSI parameter) varies. To allocate that space during execution, you can find the size of the high-level index in the DCBNCRHI field of the DCB during your DCB exit routine or after the data set is open. Use the DCBD macro instruction to gain access to the DCBNCRHI field (see "Modifying the Data Control Block" in Part 1). You can also find the size of the high-level index in the DS2NOBYT field of the index (format 2) DSCB, but you must use the utility program IEHLIST to print the information in the DSCB. You can calculate the size of the storage area required for the high-level index by using the formula

$$SMSI = \left(\begin{array}{l} \text{Number of Tracks} \\ \text{in High-Level Index} \end{array}\right)\left(\begin{array}{l} \text{Number of Entries} \\ \text{per Track} \end{array}\right) \quad (\text{Key Length} + 10)$$

The formula for calculating the number of tracks in the high-level index is in the section "Calculating Space Requirements for an Indexed Sequential Data Set" in Part 3. When a data set is shared and has the DCB integrity feature (DISP=SHR), the high-level index in storage is not updated when DCB fields are changed.

## Controlling an Indexed Sequential Data Set Device

An indexed sequential data set is processed sequentially or directly. Direct processing is accomplished by the basic access technique. Because you provide the key for the record you want read or written, all device control is handled automatically by the system. If you are processing the data set sequentially, using the queued access technique, the device is automatically positioned at the beginning of the data set.

In some cases, you may wish to process only a section or several separate sections of the data set. You do this by using the SETL macro instruction, which directs the system to begin sequential retrieval at the record having a specific key. The processing of succeeding records is the same as for normal sequential processing, except that you must recognize when the last desired record has been processed. At this point, issue the ESETL macro instruction to terminate sequential processing. You can then begin processing at another point in the data set.

### SETL—Specify Start of Sequential Retrieval

The SETL macro instruction enables you to retrieve records starting at the beginning of an indexed sequential data set or at any point in the data set. Processing that is to start at a point other than the beginning can be requested in the form of a record key, a key class (key prefix), or an actual address of a prime data record.

The key class concept is useful because you do not have to know the whole key of the first record to be processed. A key class comprises all of the keys that begin with identical characters. The key class is defined by specifying the desired characters of the key class at the address specified in the lower-limit operand of the SETL macro and setting the remaining characters to the right of the key class to binary zeros.

To use actual addresses, you must keep an account of where the records were written when the data set was created. The device address of the block containing the record just processed by a PUT-move macro instruction is available in the 8-byte data control block field DCBLPDA. For blocked records the address is the same for each record in the block.

Normally, when a data set is created with the delete option specified, deleted records cannot be retrieved using the QISAM retrieval mode. When the delete option is not specified in the DCB, the SETL macro options function as follows:

SETL B — Start at first record in the data set

SETL K — Start with record having the specified key

SETL KH — Start with record whose key is equal to or higher than the specified key

SETL KC — Start with first record having a key that falls into the specified key class

SETL I — Start with the record found at the specified direct-access address in the prime area of the data set

Because the DCBOPTCD field in the DCB can be changed after the data set is created (by respecifying the OPTCD in the DCB or DD card), it is possible to retrieve deleted records. In this case, SETL functions as noted above.

When the delete option is specified in the DCB, the SETL macro options function as follows:

SETL B — Start retrieval at first nondeleted record in the data set

SETL K — Start retrieval at record matching the specified key if that record is not deleted. If the record is deleted, an NRF (no record found) indication is set in the DCBEXCD field of the DCB, and SYNAD is given control

SETL KH — Start with first nondeleted record whose key is equal to or higher than the specified key

SETL KC — Start with first nondeleted record having a key that falls into the specified key class or follows the specified key class

SETL I — Start with first nondeleted record following the specified direct-access address

With the delete option not specified, QISAM retrieves and handles records marked for deletion like nondeleted records.

**Note:** Regardless of the SETL or delete option specified, the NRF condition will be posted in the DCBEXCD field of the DCB, and SYNAD is given control if the key or key class:

- Is higher than any key or key class in the data set

- Does not have a matching key or key class in the data set

### ESETL—End Sequential Retrieval

The ESETL macro instruction directs the system to stop retrieving records from an indexed sequential data set. A new scan limit can then be set, or processing terminated. An end-of-data-set indication automatically terminates retrieval. An ESETL macro instruction must be executed before another SETL macro instruction (described above) using the same DCB is executed.

**Note:** An ESETL macro instruction should be executed before another SETL macro instruction if the previous SETL macro instruction completed with an error.

## Creating an Indexed Sequential Data Set

You can create an indexed sequential data set in one step or in several steps. You can create the data set either by writing all records in a single step or by writing one group of records in one step and writing additional groups of records in subsequent steps. Writing records in subsequent steps is *resume loading*. When using either one step or several steps, you must present the records for writing in ascending order by key.

To create an indexed sequential data set by the one-step method, you should proceed as follows:

- Code DSORG=IS or DSORG=ISU and MACRF=PM or MACRF=PL in the DCB macro instruction.

- Specify in the DD statement the DCB attributes DSORG=IS or DSORG=ISU, record length (LRECL), blocksize (BLKSIZE), record format (RECFM), key length (KEYLEN), relative key position (RKP), options required (OPTCD), cylinder overflow (CYLOFL), and the number of tracks for a master index (NTM). Specify space requirements with the SPACE parameter. To reuse previously allocated space, omit the SPACE parameter and code DISP=(OLD, KEEP).

- Open the data set for output.

- Use the PUT macro instruction to place all the records or blocks on the direct-access volume.

- Close the data set.

The records that compose a newly created data set must be presented for writing in ascending order by key. You can merge two or more input data sets. If you want a data set with no records (a null data set), you must write at least one record when you create the data set. You can subsequently delete this record to achieve the null data set.

If records are blocked, you should not write a 1-byte record with the hexadecimal value FF. This value is used for padding; if it occurs as the last record of a block, the record cannot be retrieved.

When creating an indexed sequential data set, a procedure called *loading,* you can improve performance by using the full-track-index-write option. You do this by specifying OPTCD=U in the DCB. This causes the operating system to accumulate track-index entries in virtual storage. Note that the full-track-index-write option can be used only for fixed-length records.

If you do not specify this option, the operating system writes each normal-overflow pair of entries for the track index after the associated prime data track has been written. If you specify this option, the operating system accumulates track-index entries in virtual storage until either there are enough entries to fill a track or end-of-data or end-of-cylinder is reached. Then the operating system writes these entries as a group, writing one group for each track of track index. This option requires allocation of more storage space (the space in which the track-index entries are gathered), but the number of I/O operations required to write the index can be significantly decreased.

When you specify the full-track-index-write option, the track index entries are written as fixed-length unblocked records. If a large enough area of virtual storage is not available, the entries are written as they are created, that is, in normal-overflow pairs.

Once an indexed sequential data set has been created, its characteristics cannot be changed. However, for added flexibility, the system allows you to retrieve records using either the queued access technique with simple buffering, or the basic access technique with dynamic buffering.

**Tape-to-Disk—Indexed Sequential Data Set:** The example in Figure 59 shows the creation of an indexed sequential data set from an input tape containing 60-character records. The key by which the data set is organized is in positions 20-29. The output records will be an exact image of the input, except that the records will be blocked. One track per cylinder is to be reserved for cylinder overflow. Master indexes are to be built when the cylinder index exceeds six tracks. Reorganization information about the status of the cylinder overflow areas is to be maintained by the system. The delete option will be used during any future updating.

To create an indexed sequential data set in more than one step, create the first group of records using the one step method described above. This first section must contain at least one data record. The remaining records can then be added to the end of the data set in subsequent steps using resume load. Each group to be added must contain records with successively higher keys. This method allows you to create the indexed sequential data set in several short time periods rather than in a single long one.

This method also allows you to provide limited recovery from uncorrectable output errors. When an uncorrectable output error is detected, do not attempt to continue processing or to close the data set. If you have provided a SYNAD routine, it should issue the ABEND macro instruction to terminate processing. If no SYNAD routine is provided, the control program will terminate your processing. If the error shows that space in which to add the record was not found, you must close the data set; issuing subsequent PUT macro instructions can cause unpredictable results. You should begin recovery at the record following the end of the data as of the last successful close. The rerun time is limited to that necessary to add the new records, rather than to that necessary to recreate the whole data set.

When you extend an indexed sequential data set with resume load, the disposition parameter of the DD statement must specify MOD. To ensure that the necessary control information is in the DSCB before attempting to add records, you should at least open

```
//INDEXDD  DD       DSNAME=SLATE.DICT(PRIME),DCB=(BLKSIZE=240,CYLOFL=1,          C
//                  DSORG=IS,OPTCD=MYLR,RECFM=FB,LRECL=60,NTM=6,RKP=19,          C
                    KEYLEN=10),UNIT=3330,SPACE=(CYL,25,,CONTIG),---
//INPUTDD  DD       ---
          ...
ISLOAD    START    0
          ...
          DCBD     DSORG=IS
ISLOAD    CSECT
          OPEN     (IPDATA,,ISDATA,(OUTPUT))
NEXTREC   GET      IPDATA                          Locate mode
          LR       0,1                             Address of record in register 1
          PUT      ISDATA,(0)                      Move mode
          B        NEXTREC
          ...
CHECKERR  L        3,=A(ISDATA)                    Initialize base for errors
          USING    IHADCB,3
          TM       DCBEXCD1,X'04'
          BO       OPERR                           Uncorrectable error
          TM       DCBEXCD1,X'20'
          BO       NOSPACE                         Space not found
          TM       DCBEXCD2,X'80'
          BO       SEQCHK                          Record out of sequence
Rest of error checking
Error routine
End of job routine (EODAD FOR IPDATA)
IPDATA    DCB      ---
ISDATA    DCB      DDNAME=INDEXDD,DSORG=IS,MACRF=(PM),SYNAD=CHECKERR
          ...
```

Figure 59. Creating an Indexed Sequential Data Set

and close the data set successfully on a version of the system that includes resume load. This need be done only if the data set was created on a previous version of the system. Records may be added to the data set by resume load until the space allocated for prime data in the first step has been filled.

During resume load on a data set with a partially filled track and/or a partially filled cylinder, the track index entry and/or the cylinder index entry is overlaid when the track or cylinder is filled. If resume load abnormally terminates after these index entries have been overlaid, a subsequent resume load will result in a sequence check when it adds a key that is higher than the highest key at the last successful CLOSE but lower than the key in the overlaid index entry. When the SYNAD exit is taken for a sequence check, register 0 contains the address of the high key of the data set.

During resume load processing for variable-length records, the new records will be loaded on the next sequential prime data track. The end of the old data set may have a partially filled track.

## Retrieving and Updating an Indexed Sequential Data Set

### Sequential Retrieval and Update

To sequentially retrieve and update records in an indexed sequential data set:

- Code DSORG=IS or DSORG=ISU to agree with what you specified when you created the data set, and MACRF=GL, MACRF=SK, or MACRF=PU in the DCB macro instruction.
- Code a DD statement for retrieving the data set. The data set characteristics and options are as defined when the data set was created.
- Open the data set.
- Set the beginning of sequential retrieval (SETL).
- Retrieve records and process as required, marking records for deletion as required.
- Return records to the data set.
- Use ESETL to end sequential retrieval as required and reset the starting point.
- Close the data set to end all retrieval.

**Sequential Updates—Indexed Sequential Data Set:** Assume that, using the data set created in the previous example, you are to retrieve all records beginning with 915. Those records with a date (positions 13-16) before today's date are to be deleted. The date is in the standard form as returned by the system in response to the TIME macro instruction, that is, packed decimal 00yyddds. Overflow records can be logically deleted even though they cannot be physically deleted from the data set.

One way to solve this problem is shown in Figure 60.

### Direct Retrieval and Update

By using the basic indexed sequential access method (BISAM) to process an indexed sequential data set, you can make direct references to the records in the data set for the purpose of:

- Direct retrieval of a record by its key
- Direct update of a record
- Direct insertion of new records

```
//INDEXDD   DD          DSNAME=SLATE.DICT,---

            . . .
ISRETR      START       0
            DCBD        DSORG=IS
ISRETR      CSECT
            . . .
            USING       IHADCB,3
            LA          3,ISDATA
            OPEN        (ISDATA)
            SETL        ISDATA,KC,KEYADDR           Set scan limit
            TIME                                    Today's date in register 1
            ST          1,TODAY
NEXTREC     GET         ISDATA                      Locate mode
            CLC         19(10,1),LIMIT
            BNL         ENDJOB
            CP          12(4,1),TODAY               Compare for old date
            BNL         NEXTREC
            MVI         0(1),X'FF'                  Flag old record for deletion
            PUTX        ISDATA                      Return delete record
            B           NEXTREC
TODAY       DS          F
KEYADDR     DC          C'915'                      Key prefix
            DC          XL7'0'                      Key padding
LIMIT       DC          C'916'
            DC          XL7'0'
            . . .

CHECKERR
Test DCBEXCD1 and DCBEXDE2 for error indication
Error Routines

ENDJOB      CLOSE       (ISDATA)
            . . .
ISDATA      DCB         DDNAME=INDEXDD,DSORG=IS,MACRF=(GL,SK,PU),              C
            . . .
            SYNAD=CHECKRR
```

Figure 60. Sequentially Updating an Indexed Sequential Data Set

Because the operations are direct, there can be no anticipatory buffering. However, the system provides dynamic buffering each time a read request is made, if specified.

To ensure that the requested record is in virtual storage before you start processing, you must issue a WAIT or CHECK macro instruction. If you issue a WAIT macro instruction, you must test the exception code field of the DECB. If you issue a CHECK macro instruction, the system tests the exception code field in the DECB. If an error analysis routine has not been specified and a CHECK is issued, the program is abnormally terminated with a system completion code X'001'. In either case, if you wish to determine whether the record is an overflow record, you should test the exception code field of the DECB.

After you test the exception code field of the DECB, you need not set it to 0. If you have used a READ KU macro instruction and if you plan to use the same DECB again to rewrite the updated record using a WRITE K macro instruction, you should not set the field to 0. If you do, your record may not be rewritten properly.

To update existing records, you must use the READ KU and WRITE K combination. Because READ KU implies that the record will be rewritten in the data set, the system retains the DECB and the buffer used in the READ KU and uses them when the record is written. If you decide not to write the record, you should use the same DECB in another read or write macro instruction or issue a FREEDBUF macro instruction if dynamic buffering was used. If you issue several READ KU or WRITE K macro instructions before checking the first one, you may destroy some of your updated records unless the records are from different blocks.

If there is the possibility that your task and another task will be simultaneously accessing the same data set, or the same task has two or more DCBs opened for the same data set, you should use the DCB integrity feature. You specify the DCB integrity feature by coding DISP=SHR in your DD statement. In this way you ensure that the DCB fields are maintained for your program to process the data set correctly. If you do not use DISP=SHR and more than one DCB is open for updating the data set, the results are unpredicatable.

If you specify DISP=SHR, you must also issue an ENQ for the data set before each input/output request and a DEQ upon completion of the request. All users of the data set must use the same *qname* and *rname* operands for ENQ. For example, the users might use the data set name as the *qname* operand. For more information about using ENQ and DEQ, see *OS/VS1 Supervisor Services and Macro Instructions*.

When you are using scan mode with QISAM and you want to issue PUTX, issue an ENQ on the data set before processing it and a DEQ after processing is complete. ENQ must be issued before the SETL macro instruction, and DEQ must be issued after the ESETL macro instruction. When you are using BISAM to update the data set, do not modify any DCB fields or issue a DEQ until you have issued CHECK or WAIT.

**Sharing a BISAM DCB between Related Tasks:** When a task using BISAM processes a data set whose DCB is defined and opened by a related task, the task must issue an ENQ on the DCB before an input/output request is issued and must issue a DEQ after the WAIT or CHECK for the input/output request is issued. If the task does not enqueue the DCB and any of its related tasks terminates abnormally, the task may enter a wait state or a program check may occur. See *OS/VS1 Supervisor Services and Macro Instructions* for more information on the ENQ and DEQ macro instructions and on multitasking.

No task should process an open BISAM DCB if any other task has used it and has terminated. Otherwise, storage used by IOBs or other data areas may be freed by one task which another task might try to reuse. WAITs or CHECKs may still be issued after any of the other tasks has completed.

**Direct Update With Exclusive Control—Indexed Sequential Data Set:** In the example shown in Figure 61, the previously described data set is to be updated directly with transaction records on tape. The input tape records are 30 characters long, the key is in positions 1-10, and the update information is in positions 11-30. The update information replaces data in positions 31-50 of the indexed sequential data record.

Exclusive control of the data set is requested since more than one task may be referring to the data set at the same time. Notice that exclusive control is released after each block is written to avoid tying up the data set until the update is completed.

Note the use of the FREEDBUF macro instruction in Figure 61. Usually, the FREEDBUF macro instruction has two functions:

- To indicate to the ISAM routines that a record that has been read for update will not be written back

- To free a dynamically obtained buffer

In Figure 61, since the read operation was unsuccessful, the FREEDBUF macro instruction frees only the dynamically obtained buffer.

The first function of FREEDBUF allows you to read a record for update and then decide not to update it without performing a WRITE for update. You can use this function even when your READ macro instruction does not specify dynamic buffering, provided that you have included S (for dynamic buffering) in the MACRF field of your READ DCB.

```
//INDEXDD   DD            DSNAME=SLATE.DICT,DCB=(DSORG=IS,BUFNO=1,...),---
//TAPEDD    DD            ---
            ...
ISUPDATE    START         0
            ...
NEXTREC     GET           TPDATA,KEY
            ENQ           (RESOURCE,ELEMENT,E,,SYSTEM)
            READ          DECBRW,KU,,'S',MF=E
            WAIT          ECB=DECBRW
            TM            DECBRW+24,X'FD'                Test for any condition
            BM            RDCHECK                        but overflow
            L             3,DECBRW+16                    Pick up pointer to record
            MVC           30(20,3),UPDATE                Update record
            WRITE         DECBRW,K,MF=E
            WAIT          ECB=DECBRW
            TM            DECBRW+24,X'FD'                Any errors?
            BM            WRCHECK
            DEQ           (RESOURCE,ELEMENT,SYSTEM)
            B             NEXTREC
RDCHECK     TM            DECBRW+24,X'80'                No record found
            BZ            ERROR                          If not, go to error routine
            FREEDBUF      DECBRW,K,ISDATA                Otherwise, free buffer
            MVC           AREA+19(L'KEY),KEY             Key placed in positions 20-29
            MVC           AREA+30(L'UPDATE),UPDATE       Updated information placed in
*                                                       positions 31-50 of record
            WRITE         DECBRW,KN,,AREA-16,'S',MF=E    Add record to file
            WAIT          ECB=DECBRW
            TM            DECBRW+24,X'FD'                Test for errors
            BM            ERROR
            DEQ           (RESOURCE,ELEMENT,,SYSTEM)     Release exclusive control
            B             NEXTREC
            DS            4F                             BISAM WRITE KN work field
AREA        DS            30C                            Logical record to be added
KEY         DS            CL10
UPDATE      DS            CL20
RESOURCE    DC            CL8'SLATE'
ELEMENT     DC            C'DICT'
            READ          DECBRW,KU,ISDATA,'S','S',KEY,MF=L
ISDATA      DCB           DDNAME=INDEXDD,DSORG=IS,MACRF=(RUS,WUA),                      C
                      MSHI=INDEX,SMSI=2000
TPDATA      DCB           ---
INDEX       DS            2000C
            ...
```

Figure 61. Directly Updating an Indexed Sequential Data Set

You can effect an automatic FREEDBUF simply by reusing the DECB, that is, by issuing another READ or a WRITE KN to the same DECB. You should use this feature whenever possible, since it is more efficient than FREEDBUF. For example, in Figure 61, the FREEDBUF macro instruction could be eliminated, since the WRITE KN addressed the same DECB as the READ KU.

For an indexed sequential data set with variable-length records, you may make three types of updates by using the basic access technique. You may read a record and write it back with no change in its length, simply updating some part of the record. You do this with a READ KU followed by a WRITE K, the same way you update fixed-length records. Two other methods for updating variable-length records use the WRITE KN macro instruction and allow you to change the record length.

In one method, a record read for update (by a READ KU) may be updated in a manner that will change the record length and then be written back with its new length by a WRITE KN. In the second method, you may replace a record with another record having

the same key and possibly a different length using the WRITE KN macro instruction. To replace a record, it is not necessary to have first read the record.

In either method, when changing the record length, you must place the new length in the DECBLGTH field of the DECB before issuing the WRITE KN macro instruction. If you use a WRITE KN macro instruction to update a variable-length record that has been marked for deletion, the first bit (no record found) of the exceptional condition code field (DECBEXC1) of the DECB is set on. If this condition is found, the record must be written using a WRITE KN with nothing specified in the DECBLGTH field.

Do not try to use the DECBLGTH field to determine the length of a record read, because DECBLGTH is for use with writing records, not reading them. If you are reading fixed-length records, the length of the record read is in DCBLRECL, and if you are reading variable-length records, the length is in the record descriptor word (RDW).

**Direct Update—Indexed Sequential Data Set with Variable-Length Records:** In Figure 62, an indexed sequential data set with variable-length records is updated directly with transaction records on tape. The transaction records are of variable length and each contains a code identifying the type of transaction. Transaction code 1 indicates that an existing record is to be replaced by one with the same key; 2 indicates that the record is to be updated by appending additional information, thus changing the record length; 3 or greater indicates that the record is to be updated with no change to its length. For this example, the maximum record length of both data sets is 256 bytes. The key is in positions 6-15 of the records in both data sets. The transaction code is in position 5 of records on the transaction tape. The work area (REPLAREA) size is equal to the maximum record length plus 16 bytes.

```
//INDEXDD  DD         DSNAME=SLATE.DICT,DCB=(DSORG=IS,BUFNO=1,...),---
//TAPEDD   DD         ---
           ...
ISUPDVLR   START      0
           ...
NEXTREC    GET        TPDATA,TRANAREA
           CLI        TRANCODE,2                   Determine if replacement or
*                                                  other transaction
           BL         REPLACE                      Branch if replacement
           READ       DECBRW,KU,,'S','S',MF=E      Read record for update
           CHECK      DECBRW,DSORG=IS              Check exceptional conditions
           CLI        TRANCODE,2                   Determine if change or append
           BH         CHANGE                       Branch if change
           ...
           ...
* CODE TO MOVE RECORD INTO REPLACEA+16 AND APPEND DATA FROM TRANSACTION
* RECORD
           ...
           MVC        DECBRW+6(2),REPLAREA+16      Move new length from RDW
*                                                  into DECBLGTH (DECB+6)
           WRITE      DECBRW,KN,,REPLAREA,MF=E     Rewrite record with changed
*                                                  length
           CHECK      DECBRW,DSORG=IS
           B          NEXTREC
CHANGE     ...
           ...
* CODE TO CHANGE FIELDS OR UPDATE FIELDS OF THE RECORD
           ...
           WRITE      DECBRW,K,MF=E                Rewrite record with no
*                                                  change of length
           CHECK      DECBRW,DSORG=IS
           B          NEXTREC
REPLACE    MVC        DECBRW+6(2),TRANAREA         Move new length from RDW
*                                                  into DECBLGTH (DECB+6)
           WRITE      DECBRW,KN,,TRANAREA-16,MF=E  Write transaction record
*                                                  as replacement for record
*                                                  with the same key
           CHECK      DECBRW,DSORG=IS
           B          NEXTREC
CHECKERR   ...                                  SYNAD routine
           ...
REPLAREA   DS         CL272
TRANAREA   DS         CL4
TRANCODE   DS         CL1
KEY        DS         CL10
TRANDATA   DS         CL241
           READ       DECBRW,KU,ISDATA,'S','S',KEY,MF=L
ISDATA     DCB        DDNAME=INDEXDD,DSORG=IS,MACRF=(RUSC,WUAC),SYNAD=CHECKERR
TPDATA     DCB        ---
           ...
```

Figure 62. Directly Updating an Indexed Sequential Data Set with Variable-Length Records

# Processing a Direct Data Set

In a direct data set, there is a relationship between a control number or identification of each record and its location on the direct-access volume. This relationship allows you to gain access to a record without an index search. You determine the actual organization of the data set. If the data set has been carefully organized, location of a particular record takes less time than with an indexed sequential data set.

Although you can process a direct data set sequentially using either the queued access technique or the basic access technique, you cannot read record keys using the queued access technique. When you use the basic access technique, each unit of data transmitted between virtual storage and an I/O device is regarded by the system as a record. If, in fact, it is a block, you must perform any blocking or deblocking required. For that reason, the LRECL field is not used when processing a direct data set. Only BLKSIZE must be specified when you add or update records on a direct data set.

If dynamic buffering is specified for your direct data set, the system will provide a buffer for your records. If dynamic buffering is not specified, you must provide a buffer for the system to use.

As indicated in the discussion of direct-access devices, record keys are optional. If they are specified, they must be used for every record and must be of a fixed length.

## Organizing a Direct Data Set

In developing the organization of your data set, you can use *direct addressing*. When direct addresses are used, the location of each record in the data set is known.

If format-F records with keys are being written, the key of each record can be used to identify the record. For example, a data set with keys ranging from 0 to 4999 should be allocated space for 5000 records. Each key relates directly to a location that you can refer to as a relative record number. Therefore, each record should be assigned a unique key. If identical keys are used it is possible, during periods of high CPU and channel activity, to skip the desired record and retrieve the next record on the track. The main disadvantage of this type of organization is that records may not exist for many of the keys even though space has been reserved for them.

Space could be allocated on the basis of the number of records in the data set rather than on the range of keys. This type of organization requires the use of a cross-reference table. When a record is written in the data set, you must note the physical location either as an actual address or as a relative track and record number. The addresses must then be stored in a table that is searched when a record is to be retrieved. Disadvantages are that cross-referencing can be used efficiently only with a small data set, storage is required for the table, and processing time is required for searching and updating the table.

A more common, but somewhat complex, technique for organizing the data set involves the use of indirect addressing. In indirect addressing, the address of each record in the data set is determined by a mathematical manipulation of the key. This manipulation is referred to as randomizing or conversion. Since a number of randomizing procedures could be used, no attempt is made here to describe or explain those that might be most appropriate for your data set.

## Referring to a Record in a Direct Data Set

Once you have determined how your data set is to be organized, you must consider how the individual records will be referred to when the data set is updated or new records are added. This is important for determining whether a return address will be required when the data is created and, if so, in what form the return address will be used. The record identification can be represented in any of the following forms:

**Relative Block Address:** You specify the relative location of the record (block) within the data set as a 3-byte binary number. This type of reference can be used only with format-F records. The system computes the actual track and record number. The relative block address of the first block is 0.

**Relative Track Address:** You specify the relative track as a 2-byte binary number and the actual record number on that track as a 1-byte binary number. The relative track address of the first track is 0.

**Relative Track or Block Address and Actual Key:** In addition to the relative track or block address, you specify the address of a virtual-storage location containing the record key. The system computes the actual track address and searches for the record with the correct key.

**Actual Address:** You supply the actual address in the standard 8-byte form—MBBCCHHR. Remember that the use of an actual address may force you to indicate that the data set is unmovable.

**Extended Search:** You request that the system begin its search with a specified starting location and continue for a certain number of records or tracks. This same option can be used to request a search for unused space in which a record can be added.

To use the extended search option, you must indicate in the DCB the number of tracks (including the starting track) or records (including the starting record) that are to be searched. If you indicate a number of records, the system may actually examine more than this number. In searching a track, the system searches the whole track (starting with the first record); it therefore may examine records that precede the starting record or follow the ending record.

If the DCB specifies a number equal to or greater than the number of tracks allocated to the data set or the number of records within the data set, the entire data set is searched in the attempt to satisfy your request.

**Exclusive Control for Updating:** When more than one task is referring to the same data set, exclusive control of the block being updated is required to prevent simultaneous reference to the same record. Rather than issuing an ENQ macro instruction each time you update a block, you can request exclusive control through the MACRF field of the DCB and the type operand of the READ macro. The coding example in Figure 65 illustrates the use of exclusive control. After the READ macro instruction is executed, your task has exclusive control of the block being updated. No other task in the system requesting access to the block is given access until the operation started by your WRITE macro is complete. If, however, the block is not to be written, you can release exclusive control using the RELEX macro instruction.

**Feedback Option:** This option specifies that the system provide the address of the record requested by a READ or WRITE macro instruction. This address may be in the same form that was presented to the system in the READ or WRITE macro instruction, or as an 8-byte actual address. This option can be specified in the OPTCD parameter of the DCB and in the READ or WRITE macro instruction. If this option is omitted from the DCB but is requested in a READ or WRITE macro instruction, an 8-byte actual address is returned to the user.

The feedback option is automatically provided for a READ macro instruction requesting exclusive control for updating. This feedback will be in the form of an actual address (MBBCCHHR) unless feedback was specified in the OPTCD field of the DCB. In this case, feedback is returned in the format of the addressing scheme used in the problem program (an actual or a relative address). When a WRITE or RELEX macro instruction is issued (which releases the exclusive control that was gotten for the READ request), the system will assume that the addressing scheme used for the WRITE or RELEX macro instruction is in the same format as the addressing scheme used for feedback in the READ macro instruction.

## Creating a Direct Data Set

Once the organization of a direct data set has been determined, the process of creating it is almost identical to that of creating a sequential data set. The BSAM DCB macro instruction should be used with the WRITE macro instruction (the form used to create a direct data set). The following parameters must be specified in the DCB macro instruction:

- DSORG=PS or PSU

- DEVD=DA or omitted

- MACRF=WL

The DD statement must indicate direct-access (DSORG=DA or DAU). If keys are used, a key length (KEYLEN) must also be specified. Record length (LRECL) need not be specified but may be used to provide compatibility with sequential access method processing of this data set.

It is possible to create a direct data set using QSAM (no keys allowed) or BSAM (with or without keys and the DCB specifies MACRF=W). However, this method is not recommended because when you access this direct data set, you cannot request a function which requires the information in the capacity record (R0) data field. For example, the following restrictions would apply:

- Variable-length, undefined-length, or variable-length spanned record processing is not allowed.

- The WRITE add function with extended search for fixed-length records (with or without track overflow) is not allowed.

If a direct data set is created and updated or read within the same job step, and the OPTCD parameter is used in the creation, updating, or reading of the data set, different DCBs and DD statements should be used.

If you are using direct addressing with keys, you can reserve space for future format-F records by writing a dummy record. To reserve or truncate a track for format-U or format-V records, write a capacity record. The capacity record (R0) contains a 7-byte data field (CCHHRLL) where CCHHR is the ID of the last record on the track, and LL is the number of unused bytes on the track. If a WRITE SZ macro is issued for a track with no records, R is zero and LL is the entire length of the track.

Format-F records are written sequentially as they are presented. When a track is filled, the system automatically writes the capacity record and advances to the next track. Because of the form in which relative track addresses are recorded, direct data sets whose records are to be identified by means other than actual address must be limited in size to no more than 65,536 tracks for the entire data set.

**Tape-to-Disk—Direct Data Set:** In the example problem in Figure 63, a tape containing 204-byte records arranged in key sequence is used to create a direct data set. A 4-byte

binary key for each record ranges from 1000 to 8999, so space for 8000 records is requested.

## Adding or Updating Records on a Direct Data Set

The techniques for adding records to a direct data set depend on the format of the records and the organization used.

**Format-F With Keys:** Adding a record amounts to essentially an update by record identification. The reference to the record can be made by either a relative block address or a relative track address.

If you attempt to add a record by relative block address, the system converts the address to a relative track address. That track is searched and the new record written in place of the first dummy record on the track. If there is no dummy record on the track, you are informed that the write operation did not take place. If you request the extended search option, the new record will be written in place of the first dummy record found within the search limits you specify. If none is found, you are notified that the write operation could not take place. In the same way, a reference by relative track address causes the record to be written in place of the first dummy record on that track or the first within the search limits, if requested.

**Format-F Without Keys:** Here too, adding a record is really updating a dummy record already in the data set. The main difference is that dummy records cannot be written automatically when the data set is created. You will have to use your own method for flagging dummy records. The update form of the WRITE macro instruction

```
//DAOUTPUT DD       DSNAME=SLATE.INDEX.WORDS,DCB=(DSORG=DA,              C
//                   BLKSIZE=200,KEYLEN=4,RECFM=F),SPACE=(204,8000),---
//TAPINPUT DD        ---

         ...
DIRECT   START
         ...
         L           9,=F'1000'
         OPEN        (DALOAD,(OUTPUT),TAPEDCB)
         LA          10,COMPARE
NEXTREC  GET         TAPEDCB
         LR          2,1
COMPARE  C           9,0(2)                      Compare key of input against
*                                                control number
         BNE         DUMMY
         WRITE       DECB1,SF,DALOAD,(2)         Write data record
         CHECK       DECB1
         AH          9,=H'1'
         B           NEXTREC
DUMMY    C           9,=F'8999'                  Have 8000 records been written?
         BH          ENDJOB
         WRITE       DECB2,SD,DALOAD,DUMAREA     Write dummy
         CHECK       DECB2
         AH          9,=H'1'
         BR          10
INPUTEND LA          10,DUMMY
         BR          10
ENDJOB   CLOSE       (TAPEDCB,,DALOAD)
         ...
DUMAREA  DS          8F
DALOAD   DCB         DSORG=PS,MACRF=(WL),DDNAME=DAOUTPUT,                 C
                     DEVD=DA,SYNAD=CHECKER,---
TAPEDCB  DCB         EODAD=INPUTEND,MACRF=(GL), ---
         ...
```

Figure 63. Creating a Direct Data Set

(MACRF=W) must be used rather than the add form (MACRF=WA).

You will have to retrieve the record first (using a READ macro instruction), test for a dummy record, update, and write.

**Format-V or Format-U With Keys:** The technique used to add records in this case depends on whether records are located by indirect addressing or a cross-reference table. If indirect addressing is used, you must at least initialize each track (write a capacity record) even if no data is actually written. That way the capacity record indicates how much space is available on the track. If a cross-reference table is used, you should exhaust the input and then initialize enough succeeding tracks to contain any additions that might be required.

To add a new record, use a relative track address. The system examines the capacity record to see if there is room on the track. If there is, the new record is written. Under the extended search option, the record is written in the first available area within the search limit.

**Format-V or Format-U Without Keys:** Because a record of this type does not have a key, you can refer to the record only by its relative track or actual address (direct addressing only). When you add a record to this data set, you must retain the relative track or actual address data (for example, by updating your cross-reference table). The extended search option is not allowed because this option requires keys.

**Tape-to-Disk Add—Direct Data Set:** The example in Figure 64 involves adding records to the data set created in the last example. Notice that the write operation adds the key and the data record to the data set. If the existing record is not a dummy record, an indication is returned in the exception code of the DECB. For that reason, it is better to use the WAIT macro instruction instead of the CHECK macro instruction to test for errors or exceptional conditions.

**Tape-to-Disk Update—Direct Data Set:** The example in Figure 65 is similar to that in Figure 64, but involves updating rather than adding. There is no check for dummy records. The existing direct data set contains 25,000 records whose 5-byte keys range from 00001 to 25000. Each data record is 100 bytes long. The first 30 characters are to

```
//DIRADD    DD          DSNAME=SLATE.INDEX.WORDS,---
//TAPEDD    DD          ---
            . . .
 DIRECTAD   START
            . . .
            OPEN        (DIRECT,(OUTPUT),TAPEIN)
 NEXTREC    GET         TAPEIN,KEY
            L           4,KEY                          Set up relative record number
            SH          4,=H'1000'
            ST          4,REF
            WRITE       DECB,DA,DIRECT,DATA,'S',KEY,REF+1
            WAIT        ECB=DECB
            CLC         DECB+1(2),=X'0000'             Check for any errors
            BE          NEXTREC

            Check error bits and take required action

 DIRECT     DCB         DDNAME=DIRADD,DSORG=DA,RECFM=F,KEYLEN=4,BLKSIZE=200,      C
                        MACRF=(WA)
 TAPEIN     DCB         ---
 KEY        DS          F
 DATA       DS          CL200
 REF        DS          F
            . . .
```

Figure 64. Adding Records to a Direct Data Set

be updated. Each input tape record consists of a 5-byte key and a 30-byte data area. Notice that only data is brought into virtual storage for updating.

**Consideration for User Labels:** User labels, if desired, must be created when the data set is created. They may be updated, but not added or deleted, during processing of a direct data set. When creating a multivolume direct data set using BSAM, you should turn off the header exit entry after OPEN and turn on the trailer label exit entry just before issuing the CLOSE. This eliminates the end-of-volume exits. The first volume, containing the user label track, must be mounted when the data set is closed. If you have requested exclusive control, OPEN and CLOSE will ENQ and DEQ to prevent simultaneous reference to user labels.

**Consideration for using the 2305 Fixed Head Storage:** When a data set on a 2305 device is to be used by several tasks simultaneously, or when overlapping I/O (successive WRITEs issued without an intervening CHECK or WAIT) is used, the following combination may produce overlaying of records:

- WRITE-add processing
- Fixed records with or without track overflow

```
//DIRECTDD DD       DSNAME=SLATE.INDEX.WORDS,---
//TAPINPUT DD       ---
         ...
DIRUPDAT START
         ...
         OPEN      ( DIRECT,( UPDAT ),TAPEDCB )
NEXTREC  GET       TAPEDCB,KEY
         PACK      KEY,KEY
         CVB       3,KEYFIELD
         SH        3,=H'1'
         ST        3,REF
         READ      DECBRD,DIX,DIRECT,'S','S',0,REF+1
         CHECK     DECBRD
         L         3,DECBRD+12
         MVC       0( 30,3 ),DATA
         ST        3,DECBWR+12
         WRITE     DECBWR,DIX,DIRECT,'S','S',0,REF+1
         CHECK     DECBWR
         B         NEXTREC
         ...
KEYFIELD DS        0D'
         DC        XL3'0'
KEY      DS        CL5
DATA     DS        CL30 '
REF      DS        F
DIRECT   DCB       DSORG=DA,DDNAME=DIRECTDD,MACRF=( RISXC,WIC ),
                   OPTCD=RF,BUFNO=1,BUFL=100
TAPEDCB  DCB       ---
         ...
```

Figure 65. Updating a Direct Data Set

# PART 3: DATA SET DISPOSITION AND SPACE ALLOCATION

## Allocating Space on Direct-Access Volumes

When direct-access storage space is required for a data set, you specify the amount of space needed and the device type, and the operating system selects the device and allocates the space accordingly. This arrangement provides for flexible and efficient use of devices and available storage space, and relieves you of considering the details involved in efficient space control.

Before a direct-access volume can be used for data storage, it must be initialized by either of the utility programs IBCDASDI or IEHDASDR. The utilities' functions include in part:

- Creating the standard 80-byte volume label and writing it on cylinder 0, track 0, of the volume.

- Initializing the volume table of contents (VTOC). The location of the VTOC depends on the conventions your installation uses in initializing the volume.

- Writing the home address (HA) and capacity record (R0) for each track.

- Checking tracks and making alternate track assignments if necessary.

When the data set is to be stored on a direct-access volume, you must supply, in the DD statement, control information designating the amount of space to be allocated and the manner in which it is to be allocated.

**Note:** IEHDASDR and IBCDASDI cannot be used for an MSS 3330 virtual volume. The Access Method Services utility, CREATEV, must be used. See *OS/VS Mass Storage System (MSS) Services for Space Management* for a description of the CREATEV command.

### *Specifying Space Requirements*

The amount of space required can be specified in blocks, tracks, or cylinders. If you want to maintain device independence, specify your space requirements in blocks. If your request is in tracks or cylinders, you must be aware of such device considerations as cylinder and track capacity.

Cylinder allocation allows faster input/output of sequential data sets than does track allocation. Track allocation stops input/output at the end of every track to prevent references on the same cylinder outside of the data set. The time difference occurs when you use the sequential access method or the partitioned access method to read a data set whose record format is not fixed standard (FS). If the data set is partitioned, the time difference occurs during both loading of a module from the data set and reading of the data set's directory.

**Allocation by Blocks:** When the amount of space required is expressed in blocks, you must specify the number and average length of the blocks within the data set, as in this example:

```
//  DD SPACE=(300,(5000,100)), . . .
```

    300 = average block length in bytes
   5000 = primary quantity (number of blocks)
    100 = secondary quantity, to be allocated if the primary quantity is not sufficient (in blocks)

Note that when average block length and secondary space allocation are being used, the BLKSIZE parameter specified must be equal to the maximum block length.

From this information, the operating system estimates and allocates the number of tracks required. Space is always in whole tracks. You may also request that the space allocated for a specific number of blocks begin and end on cylinder boundaries.

You must be certain that both the quantity and the increment are large enough to contain the largest block to be written. Otherwise, all of the space requested is allocated but erased as the system tries to find a space large enough for the record.

**Allocation by Tracks or Cylinders:** The amount of space required can be expressed in tracks or cylinders, as in these examples:

```
//   DD SPACE=(TRK,(100,5)), . . .
//   DD SPACE=(CYL,(3,1)), . . .
```

**Allocation by Absolute Address:** If the data set contains location-dependent information in the form of an absolute track address (MBBCCHHR), space should be requested with respect to the number of tracks and the beginning address, as in this example:

```
//   DD SPACE=(ABSTR,(500,20)),UNIT=2314, . . .
```

where 500 tracks are required, beginning at relative track 20, which is cylinder 1, track 0.

**Allocation of Mass Storage System (MSS) Virtual Volumes:** When the data set is to be stored on an MSS virtual volume, a volume group (MSVGP) parameter may be specified instead of using the SPACE parameter on the DD card. Before the MSVGP parameter can be used, the volume group must be identified to MSS by the utility program IDCAMS.

Allocation of MSS virtual volume space should be in multiples of cylinders with secondary allocation a multiple of the primary to ensure maximum space usage and minimum fragmentation.

**Additional Space Allocation Options:** The DD statement provides you with a great deal of flexibility in specifying space requirements. These options are described in detail in *OS/VS1 JCL Reference.*

## *Estimating Space Requirements*

To determine how much space your data set requires, you must consider these variables for the device type:

- Track capacity

- Tracks per cylinder

- Cylinders per volume

- Data length (blocksize)

- Key length

- Device overhead

Figure 66 lists the physical characteristics of a number of direct-access storage devices.

The term *device overhead* refers to the space required on each track for hardware data, that is, address markers, count areas, gaps between records, record 0, etc. Device overhead varies with each device and depends also on whether the blocks are written with keys. To compute the actual space required for each block including device overhead, you can use the formulas in Figure 67. Note that any fraction of a byte must be ignored. For example, if the formula gives 15.644 bytes, you must allocate 15 bytes.

| Device | Volume Type | Maximum Block size per Track[1] | Tracks per Cylinder | Number of Cylinders[2] | Total Capacity[1,2] |
|---|---|---|---|---|---|
| 2305-1 | Drum | 14136 | 8 | 48 | 5,428,224 |
| 2305-2 | Drum | 14660 | 8 | 96 | 11,258,880 |
| 2314/2319 | Disk | 7294 | 20 | 200 | 29,176,000 |
| 3330/3333[3] (Model 1) | Disk | 13030 | 19 | 404 | 100,018,280 |
| 3330/3333 (Model 11) | Disk | 13030 | 19 | 808 | 200,036,560 |
| 3340/3344[4] | Disk | 8368 | 12 | 696 (70-megabytes) | 69,889,536 |
| | | | | 348 (35-megabytes) | 34,944,768 |
| 3350 | Disk | 19069 | 30 | 555 | 317,498,850 |

[1] Capacity indicated in bytes (when R0 is used by the IBM programming system).
[2] Excluding alternate cylinders.
[3] The Mass Storage System (MSS) virtual volumes assume the characteristics of the 3330/3333, Model 1.
[4] The 3344 is functionally equivalent to the 3340 Model 70.

Figure 66. Direct-Access Storage Device Capacities

| | | Bytes Required by Each Data Block | |
|---|---|---|---|
| Device | Track Capacity | Blocks With Keys | Blocks Without Keys |
| 2305-1 | 14568[1] | 634+KL+DL | 432+DL |
| 2305-2 | 14858[1] | 289+KL+DL | 198+DL |
| 2314/2319 | 7294 | 146+(KL+DL)534/512[2] | 101+(DL)534/512[3] |
| 3330/3333[4] (Model 1 or 11) | 13165[1] | 191+KL+DL | 135+DL |
| 3340/3344 | 8535[1] | 242+KL+DL | 167+DL |
| 3350 | 19254 | 267+KL+DL | 185+DL |

DL is data length.
KL is key length.

[1] This value is different from the maximum block size per track because the formula for the last block on the track includes an overhead for this device.
[2] The formula for the last block on the track is 45+KL+DL.
[3] The formula for the last block on the track is DL.
[4] The Mass Storage System (MSS) virtual volumes assume the characteristics of the 3330/3333, Model 1.

Figure 67. Direct-Access Device Overhead Formulas

The formulas can be combined in the following way:

If you intend to specify your space requirements in tracks (TRK) or cylinders (CYL), your estimate should be made as shown above. If you request absolute tracks (ABSTR), remember that you cannot allocate track 0, cylinder 0. The amount of space required for the VTOC will reduce the space available on the rest of the volume.

If you specify your space requirements in average block length, the system performs the computations for you.

Because a sequential data set and a direct data set are created in the same way, the estimate and specification of space requirements are identical. If you use the WRITE SZ macro instruction, your secondary allocation for a direct data set should be at least 2 tracks. Space allocation for a partitioned data set requires that you also consider the space used for the directory. Similarly, allocation for an indexed sequential data set

requires that you consider the space needed for the prime area, index areas, and overflow areas.

## *Allocating Space for a Partitioned Data Set*

What is the average size of the members to be stored on your direct-access volume? How many members will fit on the volume? Will you need directory entries for the member names only or will aliases be used? How many? Will members be added or replaced frequently? All of these questions must be answered if you are to estimate your space requirements accurately and use the space efficiently. Note, too, that a partitioned data set cannot extend beyond one volume.

If your data set will be quite large, or you expect to do a lot of updating, it might be best to allocate a full volume. If it will be small or seldom subject to change, you should make your estimate as accurate as possible to avoid wasted space or wasted time used for recreating the data set.

If the average member length is close to or less than the track length, the most efficient use of the direct-access storage space may be made with a block size of 1/3 or 1/2 the track length. For load modules, the linkage editor ignores the specified maximum block size and uses the maximum block size for the device. Program fetch always ignores BLKSIZE. It may be a good practice to indicate a block length equal to track capacity, for example, BLKSIZE=7294 for a 2314 disk. You might then ask for either 100 tracks, or 5 cylinders, thus allowing for 729,400 bytes of data.

Assuming an average length of 70,000 bytes for each member, you need space for at least 10 directory entries. If each member also has an average of three aliases, space for an additional 30 directory entries is required.

Space for the directory is expressed in 256-byte blocks. Each block contains from 3 to 20 entries, depending on the length of the user data field. If you expect 40 directory entries, request at least 8 blocks. Any unused space on the last track of the directory is wasted unless there is enough space left to contain a block of the first member. Therefore, the most advisable request in this case would be for 17 blocks.

Any of the following space specifications would cause the same size allocation for a 2314 disk:

SPACE=(7294,(100,,10))

SPACE=(CYL,(5,,10))

SPACE=(TRK,(100,,10))

Although a secondary allocation increment has been omitted in these examples, it could have been supplied to provide for extension of the member area. The directory size, however, cannot be extended.

## *Allocating Space for an Indexed Sequential Data Set*

An indexed sequential data set has three areas: prime, index, and overflow. Space for these areas can be subdivided and allocated as follows:

• Prime area—If you request a prime area only, the system automatically uses a portion of that space for indexes, taking one cylinder at a time as needed. Any unused space in the last cylinder used for index will be allocated as an independent overflow area. More than one volume can be used in most cases, but all volumes must be for devices of the same device type.

• Index area—You can request that a separate area be allocated to contain your cylinder and master indexes. The index area must be contained within one volume, but

this volume can be on a device of a different type than the one that contains the prime area volume. If a separate index area is requested, you cannot catalog the data set with a DD statement.

If the total space occupied by the prime area and index area does not exceed one volume, you can request that the separate index area be embedded in the prime area (to reduce access arm movement) by indicating an index size in the SPACE parameter of the DD statement defining the prime area.

If you request space for prime and index areas only, the system automatically uses any space remaining on the last cylinder used for master and cylinder indexes for overflow, provided the index area is on a device of the same type as the prime area.

- Overflow area—Although you can request an independent overflow area, it must be contained within one volume. If no specific request for index area is made, then it will be allocated from the specified independent overflow area.

To request that a designated number of tracks on each cylinder be used for cylinder overflow records, you must use the CYLOFL parameter of the DCB macro instruction. The number of tracks that you can use on each cylinder equals the total number of tracks on the cylinder minus the number of tracks needed for track index and for prime data, that is:

Usable tracks = total tracks - (track index tracks + prime data tracks)

Note that when you create a 1-cylinder data set, ISAM reserves 1 track on the last cylinder for the end-of-file filemark.

When you request space for an indexed sequential data set, the DD statement must follow a number of conventions, as shown below and summarized in Figure 68.

- Space can be requested only in cylinders, SPACE=(CYL,(...)), or absolute tracks, SPACE=(ABSTR,(...)). If the absolute track technique is used, the designated tracks must make up a whole number of cylinders.

- Data set organization (DSORG) must be specified as indexed sequential (IS or ISU) in both the DCB macro instruction and the DCB parameter of the DD statement.

- All required volumes must be mounted when the data set is opened; that is, volume mounting cannot be deferred.

- If your prime area extends beyond one volume, you must indicate the number of units and volumes to be spanned, for example, UNIT=(2314,3),VOLUME=(,,,3).

- You can catalog the data set using the DD statement parameter DISP=(,CATLG) only if the entire data set is defined by one DD statement; that is, if you did not request a separate index or independent overflow area.

| Criteria | | | Restrictions on Unit Types and Number of Units Requested | Resulting Arrangement of Areas |
|---|---|---|---|---|
| 1. Number of DD Statements | 2. Types of DD Statements | 3. Index Size Coded? | | |
| 3 | INDEX PRIME OVFLOW | - | None | Separate index, prime, and overflow areas. |
| 2 | INDEX PRIME | - | None | Separate index and prime areas. Any partially used index cylinder is used for independent overflow if the index and prime areas are on the same type of device. |
| 2 | PRIME OVFLOW | No | None | Prime area and overflow area with an index at its end. |
| 2 | PRIME OVFLOW | Yes | The statement defining the prime area cannot request more than one unit. | Prime area and embedded index, and overflow area. |
| 1 | PRIME | No | None | Prime area with index at its end. Any partially used index cylinder is used for independent overflow. |
| 1 | PRIME | Yes | Statement cannot request more than one unit. | Prime area with embedded index area; independent overflow in remainder of partially used index cylinder |

Figure 68. Requests for Indexed Sequential Data Sets

As your data set is created, the operating system builds the track indexes in the prime data area. Unless you request a separate index area or an embedded index area, the cylinder and master indexes are built in the independent overflow area. If you did not request an independent overflow area, the cylinder and master indexes are built in the prime area.

If an error is encountered during allocation of a multivolume data set, the IEHPROGM utility program should be used to scratch the DSCBs of the data sets that were successfully allocated. The IEHLIST utility program can be used to determine whether or not part of the data set has been allocated. The IEHLIST utility program is also useful to determine whether space is available or whether identically named data sets exist before space allocation is attempted for indexed sequential data sets. These utility programs are described in *OS/VS Utilities*.

**Specifying a Prime Data Area**

To request that the system allocate space and subdivide it as required, you should code:

```
//ddname DD DSNAME=dsname,DCB=DSORG=IS,
//         SPACE=(CYL,quantity, CONTIG),UNIT=unitname,
//         DISP=( ,KEEP),---
```

You can accomplish the same type of allocation by qualifying your *dsname* with the element indication (PRIME). This element is assumed if omitted. It is required only if you request an independent index or overflow area. To request an embedded index area when an independent overflow area is specified, you must indicate DSNAME=dsname

(PRIME). To indicate the size of the embedded index, you specify
SPACE=(CYL,(quantity,,index size)).

### Specifying a Separate Index Area

To request a separate index area, other than an embedded area as described above, you must use a separate DD statement. The element name is specified as (INDEX). The space and unit designations are as required. Notice that only the first DD statement can have a data definition name. The data set name *(dsname)* must be the same.

```
//ddname DD DSNAME=dsname(INDEX),---
//       DD DSNAME=dsname(PRIME),---
```

### Specifying an Independent Overflow Area

A request for an independent overflow area is essentially the same as for a separate index area. Only the element name, OVFLOW, is changed. If you do not request a separate index area, only two DD statements are required.

```
//ddname DD DSNAME=dsname(INDEX),---
//       DD DSNAME=dsname(PRIME),---
//       DD DSNAME=dsname(OVFLOW),---
```

### Calculating Space Requirements for an Indexed Sequential Data Set

To determine the number of cylinders required for an indexed sequential data set, you must consider the number of blocks that will fit on a cylinder, the number of blocks that will be processed, and the amount of space required for indexes and overflow areas. When you make the computations, consider how much additional space is required for device overhead. Figures 66 and 67 show device capacities and overhead formulas. In the formulas that follow, the length of the last block (or only block) must include device overhead as given in Figure 67 as Bn.

Blocks     = 1 + ((Track capacity - Length of the last block)/(Length of other blocks))
per track

$$Bt = 1 + ((Ct-Bn)/Bi)$$

The following eight steps summarize calculation of space requirements for an indexed sequential data set.

### Step 1

Once you know how many records will fit on a track and the maximum number of records you expect to create, you can determine how many tracks you will need for your data.

Number of tracks required = (Maximum number of blocks/Blocks per track) + 1

ISAM load mode reserves the last prime data track for the filemark.

Example: Assume that a 200,000 record part-of-speech dictionary is stored on an IBM 3330 Disk Storage, using the 3336 disk pack, as an indexed sequential data set. Each record in the dictionary has a 12-byte key (the word itself) and an 8-byte data area containing a part-of-speech code and control information. Each block contains 50 records; LRECL=20 and BLKSIZE=1000. Using the formula from Figure 67, we find that each track will contain 10 blocks or 500 records. A total of 401 tracks will be required for the dictionary.

$$Bt = 1 + \frac{13,165 - (191 + 12 + 1000)}{191 + 12 + 1000} = 1 + \frac{11,962}{1203} = 1 + 9 = 10$$

Records per track = (10 blocks)(50 records per block) = 500

$$\text{Prime data tracks required (T)} = \frac{200,000 \text{ records}}{500 \text{ records per track}} + 1 = 401$$

**Step 2**

You will want to anticipate the number of tracks required for cylinder overflow areas. The computation is the same as for prime data tracks, but you must remember that overflow records are unblocked and a 10-byte link field is added. Remember, if you exceed the space allocated for any cylinder overflow area, an independent overflow area is required. Those records are not placed in another cylinder overflow area.

$$\text{Overflow records per track (Ot)} = 1 + \frac{\text{Track capacity - Length of last overflow record}}{\text{Length of other overflow records}}$$

$$Ot = 1 + ((Ct-Rn)/Ri)$$

Example: Approximately 5000 overflow records are expected for the data set described in step 1. Since 56 overflow records will fit on a track, 90 overflow tracks are required. This is 90 overflow tracks for 401 prime data tracks, or approximately 1 overflow track for every 4 prime data tracks. Since the 3336 disk pack has 19 tracks per cylinder, it would probably be best to allocate 4 tracks per cylinder for overflow.

$$Ot = 1 + \frac{13,165 - (191 + 12 + 20 + 10)}{191 + 12 + 20 + 10} = 1 + \frac{12,932}{233} = 1 + 55 = 56$$

$$\text{Overflow tracks required} = \frac{5000 \text{ records}}{56 \text{ records per track}} = 90$$

Overflow tracks per cylinder (Oc) = 4

**Step 3**

You will have to set aside space in the prime area for track index entries. There will be two entries (normal and overflow) for each track on a cylinder that contains prime data records. The data field of each index entry is always 10 bytes long. The key length corresponds to the key length for the prime data records. How many index entries will fit on a track?

$$\text{Index entries per track (It)} = 1 + \frac{\text{Track capacity - Length of last index entry}}{\text{Length of other index entries}}$$

$$It = 1 + ((Ct-En)/Ei)$$

Example: Again assuming use of a 3336 disk pack and records with 12-byte keys, we find that 61 index entries will fit on a track.

$$It = 1 + \frac{13,165 - (191 + 12 + 10)}{191 + (12 + 10)} = 1 + \frac{12,952}{213} = 1 + 60 = 61$$

**Step 4**

The number of tracks required for track index entries will depend on the number of tracks per cylinder and the number of track index entries per track. Any unused space on the last track of the track index can be used for any prime data records that will fit.

$$\text{Number of track index tracks per cylinder (Ic)} = \frac{2(\text{Tracks per cylinder-overflow tracks per cylinder}) + 1}{\text{Index entries per track} + 2}$$

$$Ic = (2(Tc-Oc)+1)/(It+2)$$

Note that for variable-length records or when a prime data record will not fit on the last track of the track index, the last track of the track index is not shared with prime data records. In such a case, if the remainder of the division is less than or equal to 2, drop the remainder. In all other cases, round the quotient up to the next integer.

Example: The 3336 disk pack has 19 tracks per cylinder. You can fit 61 track index entries per track. Therefore, you need less than 1 track for each cylinder:

$$Ic = \frac{2(19-4) + 1}{61 + 2} = \frac{31}{63}$$

The space remaining on the track is (1-31/63) (13,165) = 6686 bytes.

This is enough for 6 blocks of prime data records. Since the normal number of blocks per track is 10, the blocks use 6/10 of the track, and the effective value of Ic is therefore 1-6/10 = 2/5.

Note that space is required on the last track of the track index for a dummy entry to indicate the end of the track index. The dummy entry consists of an 8-byte count field, a key field the same size as the key field in the preceding entries, and a 10-byte data field.

**Step 5**

Next you have to compute the number of tracks available on each cylinder for prime data records. You cannot include tracks set aside for cylinder overflow records.

$$\text{Prime data tracks per cylinder} = \left(\begin{array}{c}\text{Tracks} \\ \text{per cylinder}\end{array}\right) - \left(\begin{array}{c}\text{Overflow tracks} \\ \text{per cylinder}\end{array}\right) - \left(\begin{array}{c}\text{Index tracks} \\ \text{per cylinder}\end{array}\right)$$

$$Pc = Tc - Oc - Ic$$

Example: If you set aside 4 cylinder overflow tracks, and you require 2/5 of a track for the track index, 14 3/5 tracks are available on each cylinder for prime data records.

$$Pc = 19 - 4 - 2/5 = 14\ 3/5$$

**Step 6**

The number of cylinders required to allocate prime space is determined by the number of prime data tracks required divided by the number of prime data tracks available on each cylinder. This area includes space for the prime data records, track indexes, and cylinder overflow records.

$$\text{Number of cylinders required} = \text{Prime data tracks required/Prime data tracks per cylinder}$$

$$C = T/Pc$$

Example: You need 401 tracks for prime data records. You can use 14-3/5 tracks per cylinder. Therefore, 28 cylinders are required for your prime area and cylinder overflow areas.

$$C = (401)/(14\ 3/5) = 27 + \approx 28$$

**Step 7**

You will need space for a cylinder index as well as track indexes. There is a cylinder index entry for each track index (for each cylinder allocated for the data set). The size of each entry is the same as the size of the track index entries; therefore, the number of entries that will fit on a track is the same as the number of track index entries. Unused space on a cylinder index track is not shared.

$$\text{Number of tracks required for cylinder index} = (\text{Track indexes} + 1)/\text{Index entries per track}$$

$$Ci = (C+1)/It$$

Example: You have 28 track indexes (from Step 6). Since 61 index entries fit on a track (from Step 3), you need 1 track for your cylinder index. The remaining space on the last track is unused.

$$Ci = (28 + 1)/61 = 29/61 = 0.475 < 1$$

Note that every time a cylinder index crosses a cylinder boundary, ISAM writes a dummy index entry that lets ISAM chain the index levels together. The addition of dummy entries can increase the number of tracks required for a given index level. To determine how many dummy entries will be required, divide the total number of tracks required by the number of tracks on a cylinder. If the remainder is 0, subtract 1 from the quotient. If

the corrected quotient is not 0, calculate the number of tracks these dummy entries require. Also consider any additional cylinder boundaries crossed by the addition of these tracks and by any track indexes starting and stopping within a cylinder.

**Step 8**

If you have a data set large enough to require master indexes, you will want to calculate the space required according to the number of tracks for master indexes (NTM parameter) you specified in the DCB macro instruction or the DD statement.

If the cylinder index exceeds the NTM specification, an entry is made in the master index for each track of the cylinder index. If the master index itself exceeds the NTM specification, a second-level master index is started. Up to three levels of master indexes are created if required.

The space requirements for the master index are computed in the same way as those for the cylinder index.

Number of tracks
required for      = (Number of cylinder index tracks + 1)/Index entries per track
master indexes

$$M_1 = (Ci+1)/It \text{ when } Ci \geq NTM$$

$$M_2 = (M_1+1)/It \text{ when } M_1 \geq NTM$$

$$M_3 = (M_2+1)/It \text{ when } M_2 \geq NTM$$

Example: Assume that your cylinder index will require 22 tracks. Since large keys are used, only 10 entries will fit on a track. Assuming that NTM was specified as 2, 3 tracks will be required for a master index, and two levels of master index will be created.

$$M_1 = (22+1)/10 = 2.3$$

Note that every time a master index crosses a cylinder boundary, ISAM writes a dummy index entry that lets ISAM chain the index levels together. The addition of dummy entries can increase the number of tracks required for a given index level. To determine how many dummy entries will be required, divide the total number of tracks required by the number of tracks on a cylinder. If the remainder is 0, subtract 1 from the quotient. If the corrected quotient is not 0, calculate the number of tracks these dummy entries require. Also consider any additional cylinder boundaries crossed by the addition of these tracks and by any track indexes starting and stopping within a cylinder.

**Summary: Indexed Sequential Space Requirement Calculations**

1. How many blocks will fit on a track?

    $Bt = 1 + ((Ct-Bn)/Bi)$

2. How many overflow records will fit on a track?

    $Ot = 1 + ((Ct-Rn)/Ri)$

3. How many index entries will fit on a track?

    $It = 1 + ((Ct-En)/Ei)$

4. How many track index tracks are needed per cylinder?

    $Ic = (2(Tc-Oc)+1)/(It+2)$

5. How many tracks on each cylinder can be used for prime data records?

    $Pc = Tc - Oc - Ic$

6. How many cylinders are needed for the prime data area?

    $C = T/Pc$

7. How many tracks are required for the cylinder index?

$$Ci = (C+1)/It$$

8. How many tracks are required for master indexes?

$$M = (Ci+1)/It$$

# Control and Disposition of Data Sets

You specify two kinds of status and disposition information for the data sets you use for your processing by coding DISP=(status,disposition) in the disposition field of the DD statement. The first kind deals with the status of the data set when you begin processing and the relationship of the data set to other job steps in your job or other jobs. The second deals with what is to be done with the data set when you have completed processing. In the latter case, you can take advantage of the catalog of the operating system.

A data set that is being used for input has a status of OLD. If it can be used by more than one job, the status should be specified as SHR. If you are going to add to the input data set, specify MOD. The system automatically positions the access mechanism after the last record when the data set is opened. A new output data set should be indicated as NEW.

Having identified the status of the data set at the beginning of your job step, you should specify how you want it disposed of at the end of processing. If the disposition is to be unchanged, you need not specify anything. The status of an existing data set remains unchanged; a new data set is deleted. The requested disposition is performed at the end of the job step. A data set to be used in a later job can be kept (KEEP) until a subsequent request is made to delete it. If the data set is to be used by more than one job step in the same job, you can specify that it is to be passed (PASS).

If you specify the CATLG disposition, the data set name is recorded in the catalog by the system and its volume is noted. An old data set can subsequently be removed from the catalog if you specify UNCATLG.

If you wish, you can specify one disposition to be performed if the job step terminates normally, and a different disposition to be performed if the job step terminates abnormally. For example, you can specify DISP=(OLD,DELETE,KEEP) if you wish to delete a data set under normal conditions, but wish to keep it if processing is abnormally terminated. For normal termination, you can specify any disposition—PASS, KEEP, DELETE, CATLG, or UNCATLG; for abnormal termination, you can specify any disposition except PASS.

## Routing Data Through the System Input and Output Streams

The job entry subsystem is a system facility that provides spooling and scheduling of input and output data streams.

Spooling includes two basic functions:

- Input streams are read from the input device and stored on an intermediate storage device in a format convenient for later processing by the system and by the user's program.

- Output streams are similarly stored on an intermediate device until a convenient time for printing or punching.

Scheduling provides the highest degree of system availability through the orderly use of system resources that are the objects of contention.

With spooling, unit record devices are used at full rated speed if enough buffers are available, and they are used only for the time needed to read, print, or punch the data. Without spooling, the device is occupied for the entire time that a job is doing other processing. Also, because data is stored instead of being transmitted directly, output can be queued in any order and scheduled by class and by priority within each class.

You enter data into the system input stream by preceding it with a DD * or DD DATA JCL statement. This is a SYSIN data set.

Your output data can be printed or punched from a SYSOUT data set, which is called the output stream. You code the SYSOUT keyword parameter in your DD statement and designate the appropriate output class. For example, SYSOUT=A requests output class A. The class-device relationship is established for each installation; and a list of devices assigned to each output class will enable you to select the appropriate one. Refer to *OS/VS1 JCL Reference* for further information on SYSIN and SYSOUT parameters.

SYSIN and SYSOUT must be BSAM or QSAM data sets and you open and close them in the same manner as any other data set processed on a unit record device (except when multiple DCBs are used to write to the same output class, the records are not interspersed.) The DCB exit routine will be entered in the usual manner if you specify it in an exit list.

When you use QSAM with fixed-length blocked records or BSAM, the DCB block size parameter does not have to be a multiple of logical record length (LRECL) if the block size is specified through the SYSOUT DD statement. Under these conditions, if block size is greater than LRECL but not a multiple of LRECL, block size is reduced to the nearest lower multiple of LRECL when the data set is opened. This feature allows a cataloged procedure to specify blocking for SYSOUT data sets, even though your LRECL is not known to the system until execution.

Therefore, the SYSOUT DD statement of the go step of a compile-load-go procedure can specify block size without block size being a multiple of LRECL. For further information, refer to *OS/VS1 JCL Reference.*

Because a SYSOUT data set is written on a direct-access device, you should omit the DEVD operand in the DCB macro instruction, or should code DEVD=DA. Because SYSIN and SYSOUT data sets are spooled on intermediate devices, you should also avoid using device dependent macro instructions (such as FEOV, CNTRL, PRTOV, BSP, or SETPRT) in processing these data sets. (See the sections, "Device Control" and "Device Independence.")

The job entry subsystem controls all blocking and deblocking of your data to optimize system operation and ignores the number of channel programs (NCP) you specify. The block size (BLKSIZE) and number of buffers (BUFNO) specified in your program have no correlation with what is actually used by the job entry subsystem. Therefore, you can select the blocking factor that best fits your application program with no effect on the spooling efficiency of the system. For QSAM applications, move mode is as efficient as locate mode.

All record formats are allowed, except that spanned records (RECFM=VS or VBS) cannot be specified for SYSIN. Also, a record format of FIXED will be supplied if it is not specified for SYSIN.

The NOTE and POINT macro instructions may be used with SYSIN data sets on VS1 systems. However, there are restrictions associated with the usage of these macros with SYSIN. (See *OS/VS1 Data Management Macro Instructions.*)

Your program is responsible for printing format, pagination, and header control. You can supply control characters for SYSOUT data sets in the normal manner by specifying ANSI or machine characters in the DCB. Standard controls are provided by default if

they are not specified. The length of output records must not exceed the allowable maximum length for the ultimate device. Cards can be punched in EBCDIC mode only.

Your SYNAD routine will be entered if an error occurs during data transmission to or from an intermediate storage device. Again, because the specific device is indeterminate, your SYNAD routine code should be device independent.

**Note:** For spooled data sets, the area address must be in the same protection key as that specified in the TCB (TCBPKF field). If not, a system 001 ABEND will occur.

## Concatenating Sequential and Partitioned Data Sets

Two or more sequential or partitioned data sets can be automatically retrieved by the system and processed successively as a single data set. This reading technique is known as *concatenation*. A maximum of 255 data sets (16, if partitioned) can be concatenated, but they must be used only for input.

To save time when processing two consecutive data sets on a single volume, you specify LEAVE in your OPEN macro instruction. Concatenated data sets cannot be read backward.

When data sets are concatenated, the system treats the group as a single data set and only one data extent block (DEB) is constructed. Thus, it is important to consider the characteristics of the individual data sets being concatenated. Data sets with like characteristics are those that may be processed correctly using the same data control block (DCB), input/output block (IOB), and channel program. Any exception makes them unlike. Concatenated partitioned data sets are always treated as like and use the attributes of the first data set only. You must inform the system if unlike data sets are concatenated by modifying the DCBOFLGS field of the DCB. The indication must be made before the end of the current data set is reached. You must set bit 4 to 1 by using the instruction OI DCBOFLGS,X'08' as described in "Modifying the Data Control Block." If bit 4 of the DCBOFLGS field is 1, end-of-volume processing for each data set will issue a CLOSE for the data set just read and an OPEN for the next concatenated data set. This opening and closing procedure updates the fields in the DCB and, if necessary, builds a new IOB and a new channel program. If the buffer pool was obtained automatically by the open routine, the procedure also frees the buffer pool and obtains a new one for the next concatenated data set. The procedure does not issue a FREEPOOL for the last concatenated data set. Unless you have some way of determining the characteristics of the next data set before it is opened, you should not reset the DCBOFLGS field to indicate like characteristics during processing.

When unlike data sets have been concatenated, you should not issue multiple input requests, that is, a series of READ or GET macro instructions, in your program. If you do, you will have to arrange some way to determine which requests have been completed and which must be reissued. In any case, the GET or READ macro instruction that detected the end of data set will have to be reissued. Figure 69 illustrates a possible routine for determining when a GET or READ must be reissued. This restriction does not apply to like data sets since no open or close operation is necessary between data sets.

When the change from one data set to another is made, label exits are taken as required; automatic volume switching is also performed for multiple-volume data sets unless they are partitioned. Your end-of-data-set (EODAD) routine is not entered until the last data set has been processed, except that for partitioned data sets, your EODAD routine receives control at the end of each member. At that time, you can process the next member or close the data set.

You process a concatenation of partitioned data sets the same way you process a single partitioned data set with one exception. You must use the FIND macro instruction to
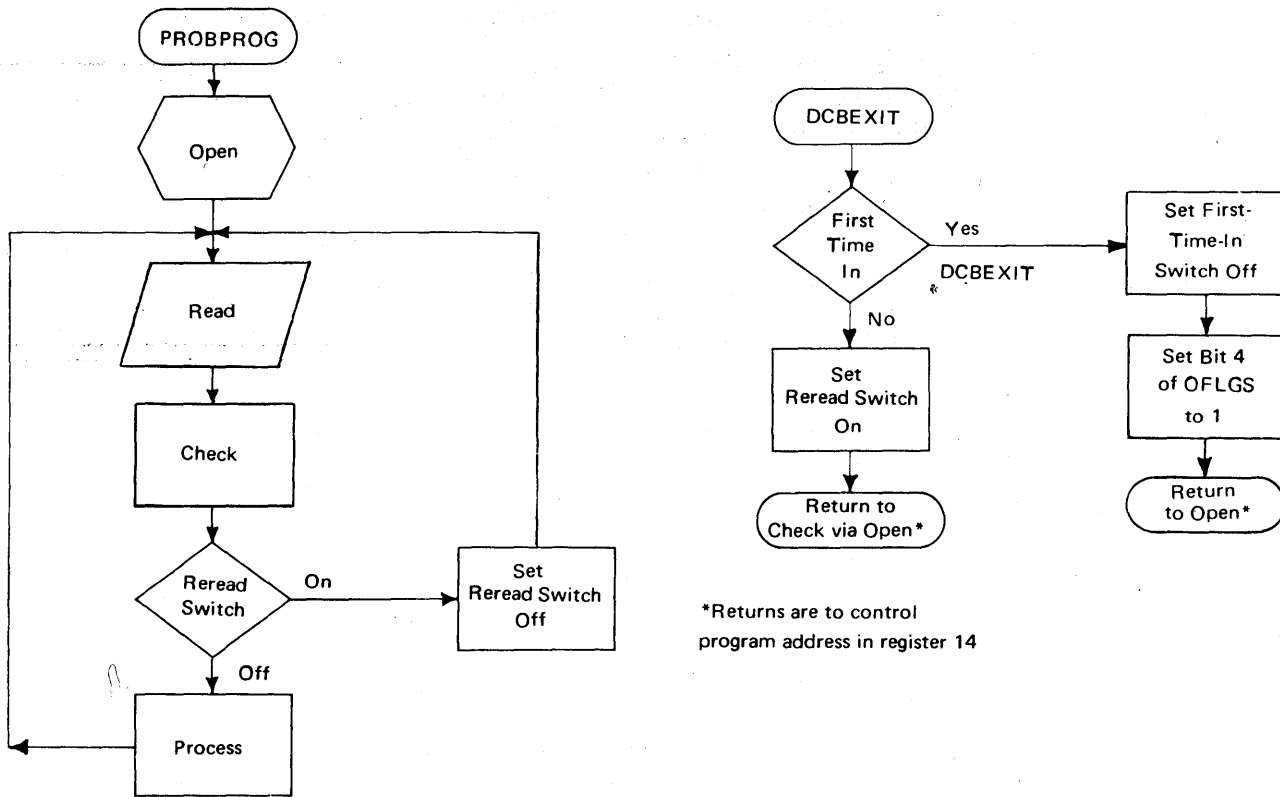
Figure 69. Reissuing a READ for Unlike Concatenated Data Sets

begin processing a member; you cannot use the POINT (or NOTE) macro instruction until after the FIND macro instruction has been issued. Figure 53 shows how to process a single partitioned data set using FIND. If two members of different data sets in the concatenation have the same name, the FIND macro instruction determines the address of the first one in the concatenation. You would not be able to process the second one in the concatenation. The BLDL macro instruction provides the concatenation number of the data set to which the member belongs in the K field of the build list. See the section "BLDL—Construct a Directory Entry List" in Part 2 of this book.

If issuing an RDJFCB macro, see the RDJFCB macro instruction in *OS/VS1 Data Management for System Programmers.*

## Rotational Position Sensing Considerations

Direct-access storage devices with the rotational position sensing (RPS) feature (for example, the 3330) usually employ channel programs that are not compatible with direct-access storage devices that lack the RPS feature. Therefore, if you concatenate otherwise "like" data sets residing on devices both with and without the RPS feature, standard (nonRPS) channel programs will be used, with a resultant loss of the I/O overlap efficiency of rotational position sensing. Concatenated partitioned data sets are always treated as "like" data sets, regardless of how the DCBOFLGS field is set in the DCB. Data sets with undefined length records and track overflow (RECFM=UT specified in the DCB) are not processed with the RPS feature.

On the other hand, if you concatenate sequential data sets with "unlike" attributes, you'll get RPS channel programs for the data sets residing on RPS devices, unless any following direct-access concatenations are nonRPS devices.

Further discussion and examples of concatenated data sets are contained in *OS/VS1 JCL Reference.*

## Cataloging Data Sets

The nonVSAM catalog is itself a data set residing on one or more direct-access volumes. It is organized into levels of indexes that connect the data-set names to corresponding volumes and data-set sequence numbers. For each level of qualification in the data-set name, there is an index group in the catalog.

The highest level of the catalog resides on the system-residence volume. The VTOC contains an entry for the DSCB defining the catalog and its highest-level index, the volume index. The lowest-level index contains the simple name of the data set and the number of the volume on which it resides.

The complete catalog can exist on the system-residence volume, or you can specify that parts of it be constructed on other volumes. Any volume containing part of the catalog is called a *control volume.* The use of control volumes allows data sets that are functionally related to be cataloged separately. The advantages include:

- Control volumes can be moved from one processing system to another.

- System-residence requirements can be reduced by placement of seldom-used indexes on a control volume.

For any given data set, only one level of control volume, other than the system-residence volume, can be used. Notice that in Figure 70, index E, which is the highest-level index on the control volume, has an entry in both volume indexes.
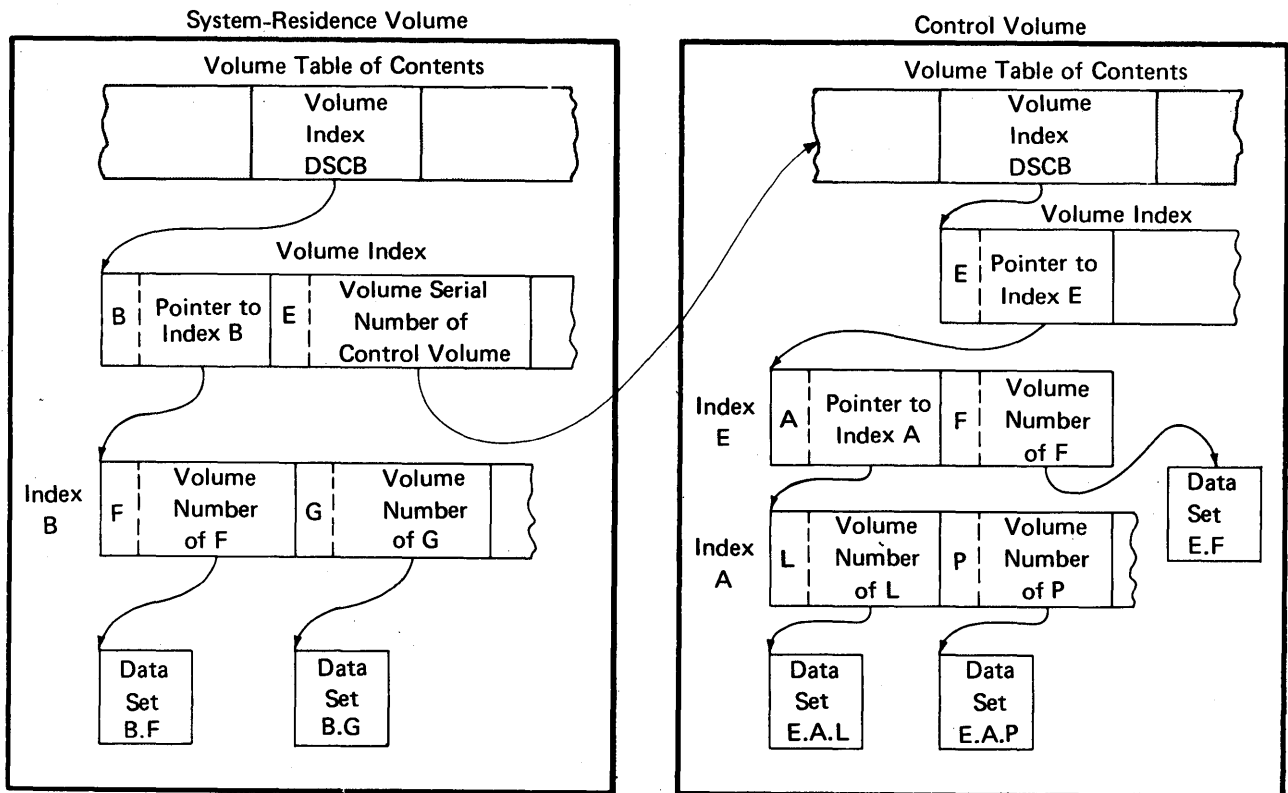


Figure 70. Catalog Structure on Two Volumes

All VSAM data sets must be cataloged in a VSAM catalog. However, nonVSAM data sets can be cataloged in a VSAM catalog instead of the system catalog. For further information on the VSAM catalog, see *OS/VS1 Access Method Services.*

Permanent Mass Storage System (MSS) data sets should be cataloged to allow efficient use of the Mass Storage Volume Control (MSVC) functions. For information on MSVC, see *OS/VS Mass Storage System (MSS) Services for Space Management.*

### Entering a Data Set Name in the Catalog

The catalog structure, including all levels of indexes, is initially created by the utility program, IEHPROGM, or when a qualified data set name is cataloged. A data set name can then be entered if the proper index levels of the name exist.

For example, if a data set named A.B.C is to be cataloged, the volume index on the system-residence volume must have an index entry for index A, which must point to an index B. When the data set A.B.C is cataloged, C is entered into index B along with the volume serial number of the volume where data set A.B.C resides. The cataloging request is entered as:

```
//ddname DD DSNAME=A.B.C,DISP=( ,CATLG )
```

The Access Method Services program may be used to catalog a data set (either VSAM or nonVSAM) in a VSAM catalog. For more information on the Access Method Services program, see *OS/VS1 Access Method Services.*

# Generation Data Groups

A generation data group is a group of related cataloged data sets. The manner in which these data sets are cataloged is what makes them a generation data group. Within a generation data group, the generations can have like or unlike DCB attributes and data set organizations. If the attributes and organizations of all generations in a group are identical, the generations can be retrieved together as a single data set. Each data set within a generation data group is called a generation data set. Generation data sets are sometimes called *generations.*

There are advantages to grouping related data sets. Because the catalog management routines can refer to the information in a special index—called a *generation index*—in the catalog:

- All of the data sets in the group can be referred to by a common name.

- The operating system is able to keep the generations in chronological order.

- Outdated or obsolete generations can be automatically deleted by the operating system.

The management of a generation data group depends upon the fact that generation data sets have sequentially ordered names—absolute and relative names—that represent their age. The absolute generation name is the representation used by the catalog management routines in the catalog. Older data sets have smaller absolute numbers. The relative name is a signed integer used to refer to the latest (0), next to the latest (−1), etc. generation. The relative number can also be used to catalog a new generation (+1).

## *Absolute Generation and Version Numbers*

An absolute generation and version number is used to identify a specific generation of a generation data group. The generation and version numbers are in the form GxxxxVyy, where xxxx is an unsigned four-digit decimal generation number and yy is an unsigned two-digit decimal version number. For example:

- A.B.C.G0000V00 is generation data set zero, version zero in the generation data group A.B.C.

- A.B.C.G0001V00 is generation data set one, version zero in generation data group A.B.C.

- A.B.C.G0009V01 is generation data set nine, version one in generation data group A.B.C.

The number of new generations and versions is limited by the number of digits in the absolute generation name, that is, 9999 for generations and 99 for versions.

The generation number is automatically maintained by the system. The number of generations kept depends on the size of the generation index. For example, if the size of the generation index allows ten entries, the ten latest generations may be maintained in the generation data group.

The version number allows you to perform normal data set operations without disrupting the management of the generation data group. For example, if you want to update the second generation in a three-generation group, replace generation two, version zero, with generation two, version one. Only one version is kept per generation.

A generation can be cataloged using either absolute or relative numbers. When a generation is cataloged, a generation and version number is placed as a low level entry in the generation data group. In order to catalog a version number other than V00, you must use an absolute generation and version number.

## *Relative Generation Number*

As an alternative to using absolute generation and version numbers when cataloging or referring to a generation, you can use a relative generation number. To specify a relative number, use the generation data group name followed by a negative integer, a positive integer, or a zero, enclosed in parentheses. For example, A.B.C(-1). A.B.C(+1), or A.B.C(0).

The value of the specified integer tells the operating system what generation number to assign to a new generation, or it tells the system the location (in the generation index) of an entry representing a previously cataloged generation.

When you use a relative generation number to catalog a generation, the operating system assigns an absolute generation number and a version number of V00 to represent that generation. The absolute generation number assigned depends on the number last assigned and the value of the relative generation number that you are now specifying. For example, if in a previous job, generation A.B.C.G0005V00 was the last generation cataloged, and you specify A.B.C(+1), the generation now cataloged is assigned the number G0006V00. Though any positive relative generation number can be used, a number greater than 1 may cause absolute generation numbers to be skipped.

When you use a relative generation number to refer to a generation that was cataloged in a previous job, the relative number has the following meaning:

- A.B.C(0) refers to the latest existing cataloged entry.

- A.B.C(-1) refers to the next-to-the-latest entry, etc.

When cataloging is requested via JCL, all actual cataloging occurs at job termination—not at step termination. Because this is so:

- A relative number used in the JCL refers to the same generation throughout a job.

- A job step that terminates abnormally may be deferred for a later step restart. If the step cataloged a generation data set via JCL, you must change all relative generation numbers in the succeeding steps via JCL before resubmitting the job.

For example, if the succeeding steps contained the relative generation numbers:

- A.B.C(+1), which refers to the entry cataloged in the terminated job step.

- A.B.C(0), which refers to the next to the latest entry.

- A.B.C(-1), which refers to the latest entry, prior to A.B.C(0).

You must change them as follows before the step can be restarted: A.B.C(0), A.B.C(-1), A.B.C(-2), etc.

Note: New generation data group requests are cataloged with a volume serial number of X'FF4040404040' if they are not opened, so that data set integrity is maintained and an incorrect generation is not retrieved.

## Building a Generation Index

A generation data group is managed via the information found in a generation index. (Note that an alias name cannot be assigned to the highest level of a generation index.) The BLDG function builds the index. The BLDG function also indicates how older or obsolete generations are to be handled when the index is full. For example, when the index is full, you may wish to empty it, scratch existing generations, and begin cataloging a new series of generations.

After the index is built, a generation can be cataloged by its generation data group name and either an absolute generation and version number or a relative generation number.

Examples on how to build a generation-data-group index are found in *OS/VS Utilities*.

## Creating a New Generation

To create a new generation data set you must first allocate space for the generation, then catalog the generation in the system catalog.

### Allocating a Generation

To take full advantage of the facilities of the system, the allocation can be patterned after a previously allocated generation in the same group. This is accomplished by the specification of DCB attributes for the new generation as described below.

If you are using absolute generation and version numbers, DCB attributes for a generation can be supplied directly in the DCB parameter of the DD statement defining the generation to be created and cataloged.

If you are using relative generation numbers to catalog generations, DCB attributes can be supplied either: (1) by creating a model DSCB on the volume on which the index resides (the volume containing the SYSCTLG data set) or (2) by referring to a cataloged data set for the use of its attributes. Attributes can be supplied before you catalog a generation, when you catalog it, or at both times, as follows:

1. Create a model DSCB on the volume on which your index resides. You can provide initial DCB attributes when you create your model; however, you need not provide any attributes at this time. Since only the attributes in the data set label are used, the

model data set should be allocated with SPACE=(TRK,0) to conserve direct-access space. Initial or overriding attributes can be supplied when you create and catalog a generation.[1] To create a model DSCB, include the following DD statement in the job step that builds the index or in any other job step that precedes the step in which you create and catalog your generation.

```
//name DD DSNAME=datagrpname,DISP=( ,KEEP),SPACE=(TRK,(0)),
//         UNIT=yyyy,VOLUME=SER=xxxxxx,
//         DCB=(applicable subparameters)
```

The DSNAME is the common name by which each generation is identified; xxxxxx is the serial number of the volume containing the catalog. If no DCB subparameters are desired initially, you need not code the DCB parameter.

2. You do not need to create a model DSCB if you can refer to a cataloged data set whose attributes are identical to those you desire or to an existing model DSCB for which you can supply overriding attributes. A cataloged data set referred to in this manner must reside on the same volume as your index. To refer to a cataloged data set for the use of its attributes, specify DCB=(dsname) on the DD statement that creates and catalogs your generation. To refer to an existing model, specify DCB=(modeldscbname, your attributes) on the DD statement that creates and catalogs your generation.

### Cataloging a Generation

A generation can be cataloged through the use of normal job control language procedures or through the use of IEHPROGM.

**Using JCL to Catalog a Generation:** Assuming that a generation data group index has been built and that provisions have been made for supplying DCB attributes, a generation is created and cataloged in the same manner as any other type of data set.

Whether you use relative numbers in job control language procedures or absolute generation and version numbers, you need not catalog the new generation immediately.

**Using IEHPROGM to Catalog a Generation:** The CATLG function of IEHPROGM can be used to catalog a generation. Again, the prerequisite for cataloging a generation is the existence of a generation data group index in the SYSCTLG data set.

**Note:** You must always use an absolute generation and version number to catalog or uncatalog a generation using IEHPROGM. (IEHMOVE and IEHLIST also require that absolute generation and version numbers be used.)

### Passing a Generation

A new generation may be passed when created. That generation may then be cataloged in a succeeding job step or deleted at the end of the job as in normal disposition processing when DISP=(,PASS) is specified on the DD statement.

However, once a generation has been created with DISP=(NEW,PASS) specified on the DD statement, another new generation for that data group must not be cataloged until the passed version has been deleted or cataloged. To do so would cause the wrong generation to be used when referencing the passed generation data set. If that data set was later cataloged, a bad generation would be cataloged and a good one lost.

For example, if A.B.C(+1) was created with DISP=(NEW,PASS) specified on the DD statement, then A.B.C.(+2) must not be created with DISP=(NEW,CATLG) until A.B.C(+1) has been cataloged or deleted.

---

[1] Only one model DSCB is necessary for any number of generations. If you plan to use only one model, do not supply DCB attributes when you create the model. When you subsequently create and catalog a generation, include necessary DCB attributes in the DD statement referring to the generation. In this manner, any number of generation data groups can refer to the same model. Note that the catalog and model data set label are always located on a direct-access volume, even for a magnetic tape generation data group.

By using the proper JCL, the advantages to this support are:

- JCL will not have to be changed in order to rerun the job.

- The lowest generation version will not be deleted from the index until a valid version is cataloged.

### Creating an ISAM Data Set as Part of a Generation Data Group

To create an indexed-sequential data set as part of a generation data group, you must:
(1) create the indexed-sequential data set separately from the generation group and
(2) use IEHPROGM to put the indexed-sequential data set into the generation group.

Use the RENAME function to rename the data set. Then use the CATLG function to catalog the data set. For instance, if MASTER is the name of the generation data group, and GggggVvv is the absolute generation name, you would code the following:

```
RENAME DSNAME=ISAM,VOL=2314=SCRTCH,NEWNAME=MASTER.GggggVvv
CATLG DSNAME=MASTER.GggggVvv,VOL=2314=SCRTCH
```

## *Retrieving a Generation*

A generation may be retrieved through the use of job control language procedures. Any operation that can be applied to a non-generation data set can be applied to a generation. For example, a generation can be updated and reentered in the catalog, or it can be copied, printed, punched, or used in the creation of new generation or non-generation data sets.

You can retrieve a generation by using either relative generation numbers or absolute generation and version numbers.

Because two or more jobs can compete for the same resource, generation data groups should be updated with care, as follows:

- No two jobs running concurrently should refer to the same generation data group. As a partial safeguard against this situation, use absolute-generation and version numbers when cataloging or retrieving a generation in a multiprogramming environment. If you use relative numbers, a job running concurrently may update the generation data group index, perhaps cataloging a new generation which you will then retrieve in place of the one you wanted.

- Even when using absolute generation and version numbers, a job running concurrently might catalog a new version of a generation or perhaps delete the generation you wished to retrieve. For this reason, some degree of control should be maintained over the execution of job steps referring to generation data groups.

# Controlling Confidential Data

## *Password Protection for NonVSAM Data Sets*

Password protection as described here applies to nonVSAM data sets only. For information on password protection for VSAM data sets, see *OS/VS1 Access Method Services.*

In addition to the usual label protection that prevents opening of a data set without the correct data set name, the operating system provides data set security options that prevent unauthorized access to confidential data. Two levels of protection options are available. You specify these options in the LABEL field of a DD statement with the parameter PASSWORD or NOPWREAD.

- Password protection (specified by the PASSWORD parameter) makes a data set unavailable for all types of processing until a correct password is entered by the system operator.

- No-password-read protection (specified by the NOPWREAD parameter) makes a data set available for input without a password, but requires that the password be entered for output or delete operations.

If an incorrect password is entered twice, the job is terminated by the system if it is being requested by the open or EOV routine. For a scratch or rename request, a return code is given.

You can request password protection when you create the data set by using the LABEL field of the DD statement in your JCL. The system sets the data set security byte either in the standard header label 1 as shown in *OS/VS Tape Labels* or in the identifier data set control block (DSCB) as shown in *OS/VS1 System Data Areas.* Once you have requested security protection for magnetic tapes, you cannot remove it with JCL unless you recreate the data set and scratch the protected data set.

In addition to requesting password protection in your JCL, you must enter at least one record for each protected data set in a data set named PASSWORD that must be created on the system-residence volume. You should also request password protection for the PASSWORD data set itself to prevent both reading and writing without knowledge of the password.

For a data set on a direct-access device you can place the data set under protection at the same time that you enter its password in the PASSWORD data set. You can use the PROTECT macro instruction or the IEHPROGM utility program to add, change, or delete an entry in the PASSWORD data set; with either of these methods the system updates the DSCB of the data set to reflect its protected status. This provision eliminates the need for you to use JCL whenever you add, change, or remove security protection for a data set on a direct-access device. *OS/VS1 Data Management for System Programmers* describes how to maintain the PASSWORD data set, including the PROTECT macro instruction; *OS/VS Utilities* describes the IEHPROGM utility program.

# APPENDIX A: DIRECT-ACCESS LABELS

Only standard label formats are used on direct-access volumes. Volume, data set, and optional user labels are used (see Figure 71). In the case of direct-access volumes, the data set label is the data set control block (DSCB).
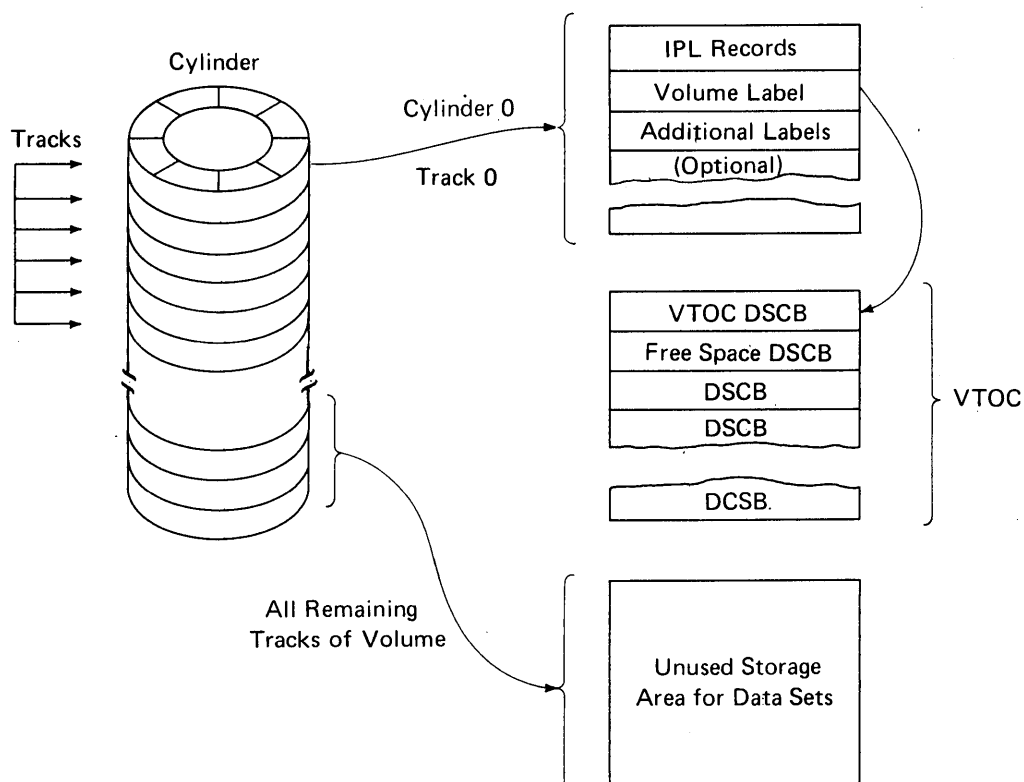


Figure 71. Direct-Access Labeling

## Volume-Label Group

The volume-label group immediately follows the first two initial program loading (IPL) records on track 0 of cylinder 0 of the volume. It consists of the initial volume label at record 3 plus a maximum of seven additional volume labels. The initial volume label identifies a volume and its owner, and is used to verify that the correct volume is mounted. It can also be used to prevent use of the volume by unauthorized programs. The additional labels can be processed by an installation routine that is incorporated into the system.

The format of the direct-access volume label group is shown in Figure 72.

(Up to 7 Additional Volume Labels)
80-Byte Physical Record

| Field | | | |
|---|---|---|---|
| 1 | (3) | Volume Label Identifier (VOL) |
| 2 | (1) | Volume Label Number (1) |
| 3 | (6) | Volume Serial Number |
| 4 | (1) | Volume Security |
| 5 | (5) | VTOC Pointer |
| 6 | (25) | Reserved (Blank) |
| 7 | (10) | Owner Name and Address Code |
| 8 | (29) | Blank |
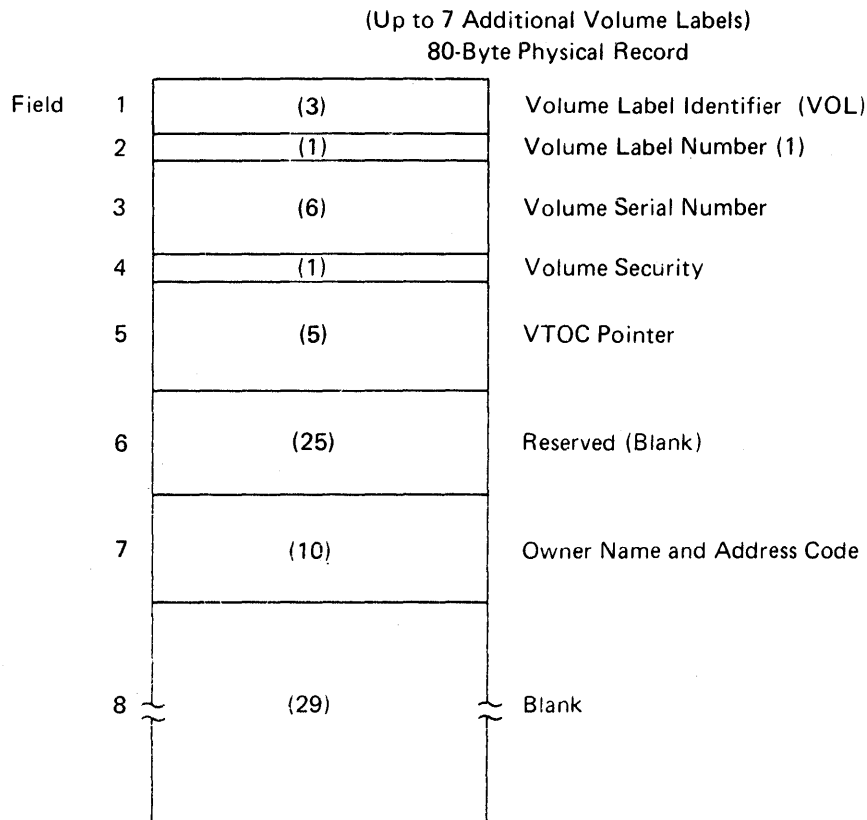
Figure 72. Initial Volume Label

## Initial Volume Label Format

The 80-byte initial volume label is preceded by a four-byte key containing VOL1.

**Volume Label Identifier (VOL):** Field 1 identifies a volume label.

**Volume Label Number (1):** Field 2 identifies the relative position of the volume label in a volume label group. It must be written as C'1'.

The operating system identifies an initial volume label when, in reading the initial record, it finds that the first 4 characters of the record are VOL1.

**Volume Serial Number:** Field 3 contains a unique identification code assigned when the volume enters the system. You can place the code on the external surface of the volume for visual identification. The code is normally numeric (000001-999999), but may be any 1 to 6 alphameric or national (#, $, @) characters, or a hyphen (X'60'). If this field is less than 6 characters, it is padded on the right with blanks.

**Volume Security:** Field 4 is reserved for use by installations that wish to provide security for volumes. Make this field a C'0' unless you have your own security processing routines.

**VTOC Pointer:** Field 5 of direct-access volume label 1 contains the address of the VTOC in the form of CCHHR.

**Reserved:** Field 6 is reserved for future developmental purposes. Leave it blank.

**Owner Name and Address Code:** Field 7 contains a unique identification of the owner of the volume.

All of the bytes in Field 8 are left blank.

## Data Set Control Block (DSCB)

The system automatically constructs a DSCB when space is requested for a data set on a direct-access volume. Each data set on a direct-access volume has one or more DSCBs to describe its characteristics. The DSCB appears in the VTOC and contains operating-system data, device-dependent information, and data set characteristics, in addition to space allocation and other control information. There are seven kinds of DSCBs, each with a different purpose and a different format number. For an explanation of the seven kinds of DSCBs, see *OS/VS1 System Data Areas*.

## User Label Groups

User header and trailer label groups can be included with data sets of physically sequential or direct organization. The labels in each group have the format shown in Figure 73.
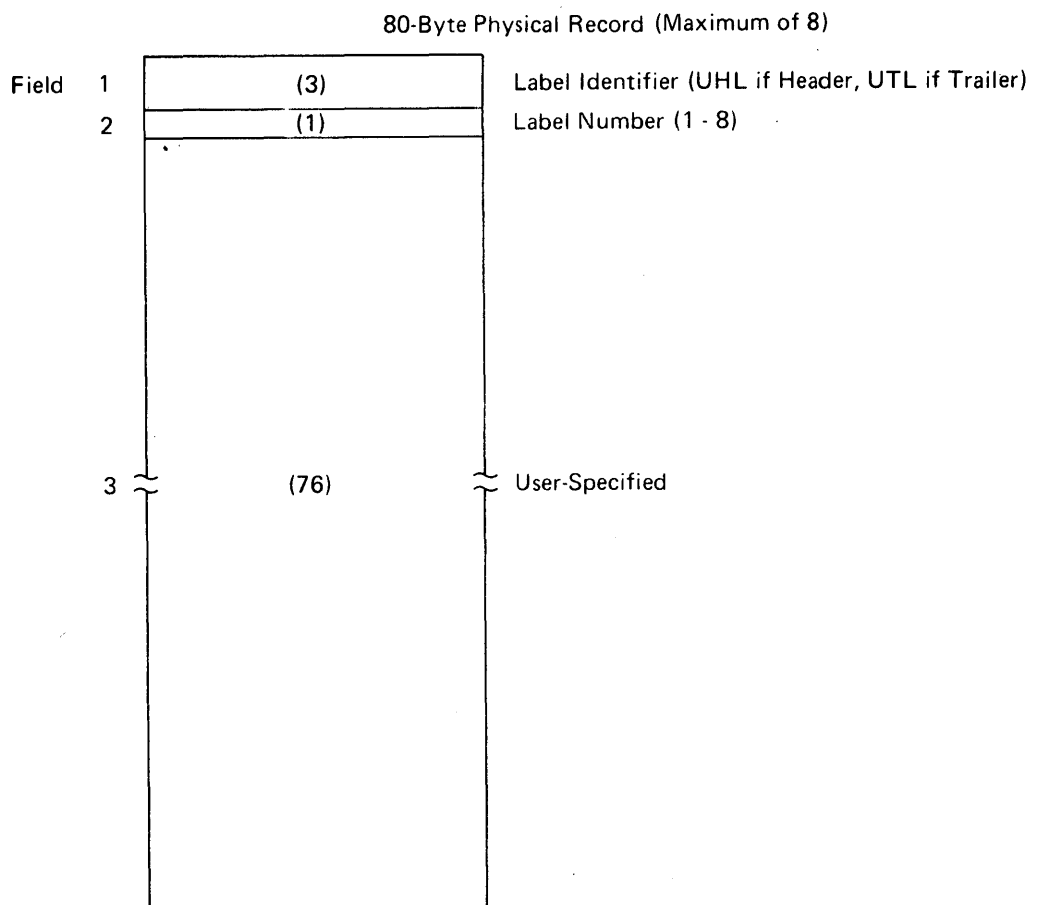


Figure 73. User Header and Trailer Labels

Each group can include up to eight labels, but the space required for both groups must not be more than 1 track on a direct-access device. The current minimum track size

allows a maximum of eight labels, including both header and trailer labels. Consequently, a program becomes device-dependent (among direct-access devices) when it creates more than eight labels.

If user labels are specified in the DD statement (LABEL=SUL), an additional track is normally allocated when the data set is created. No additional track is allocated when specific tracks are requested (SPACE=(ABSTR,...)), or when tracks allocated to another data set are requested (SUBALLOC=...). In either case, labels are written on the first track that is allocated.

User Header Label Group: The operating system writes these labels as directed by the processing program recording the data set. The first 4 characters of the user header label must be UHL1,..., UHL8; you can specify the remaining 76 characters. When the data set is read, the operating system makes the user header labels available to the problem program for processing.

User Trailer Label Group: These labels are recorded (and processed) as explained in the preceding text for user header labels, except that the first 4 characters must be UTL1,...., UTL8.

## User Header and Trailer Label Format

**Label Identifier:** Field 1 indicates the kind of user header label. UHL indicates a user header label; UTL indicates a user trailer label.

**Label Number:** Field 2 identifies the relative position (1-8) of the label within the user label group.

**User-Specified:** Field 3 (76 bytes).

# APPENDIX B: CONTROL CHARACTERS

As an optional feature, each logical record, in any record format, may include a control character. This control character is recognized and processed if a data set is being written to a printer or punch.

For format-F and format-U records, this character is the first byte of the logical record.

For format-V records, it must be the fifth byte of the logical record, immediately following the record descriptor word.

Two options are available. If either option is specified in the DCB, the character must appear in every record and other line spacing or stacker selection options also specified in the DCB are ignored.

## Machine Code

You can specify in the DCB that the machine code control character has been placed in each logical record. If the record is to be written, the appropriate byte must contain the command code bit configuration specifying both the write and the desired carriage or stacker select operation.

The machine code control characters for a printer are as follows:

| Print and Then Act | | Act Immediately (No Printing) |
|---|---|---|
| Code in Hexadecimal | Action | Code in Hexadecimal |
| 01 | Print only (no space) | |
| 09 | Space 1 line | 0B |
| 11 | Space 2 lines | 13 |
| 19 | Space 3 lines | 1B |
| 89 | Skip to channel 1 | 8B |
| 91 | Skip to channel 2 | 93 |
| 99 | Skip to channel 3 | 9B |
| A1 | Skip to channel 4 | A3 |
| A9 | Skip to channel 5 | AB |
| B1 | Skip to channel 6 | B3 |
| B9 | Skip to channel 7 | BB |
| C1 | Skip to channel 8 | C3 |
| C9 | Skip to channel 9 | CB |
| D1 | Skip to channel 10 | D3 |
| D9 | Skip to channel 11 | DB |
| E1 | Skip to channel 12 | E3 |

The machine code control characters for a card read punch device are as follows:

| Code in Hexadecimal | Action |
| --- | --- |
| 01 | Select stacker 1 |
| 41 | Select stacker 2 |
| 81 | Select stacker 3 |

Other command codes for specific devices are contained in publications describing the control units and devices.

# Extended American National Standards Institute Code

In place of machine code, you can specify control characters defined by the American National Standards Institute, Inc. (ANSI). Whenever IBM publications refer to ANSI code, they are as follows:

| Code | Action Before Printing a Line |
| --- | --- |
| b | Space one line (blank code) |
| 0 | Space two lines |
| - | Space three lines |
| + | Suppress space |
| 1 | Skip to channel 1 |
| 2 | Skip to channel 2 |
| 3 | Skip to channel 3 |
| 4 | Skip to channel 4 |
| 5 | Skip to channel 5 |
| 6 | Skip to channel 6 |
| 7 | Skip to channel 7 |
| 8 | Skip to channel 8 |
| 9 | Skip to channel 9 |
| A | Skip to channel 10 |
| B | Skip to channel 11 |
| C | Skip to channel 12 |

| Code | Action After Punching a Card |
| --- | --- |
| V | Select punch pocket 1 |
| W | Select punch pocket 2 |

These control characters include those defined by ANSI FORTRAN. If any other character is specified, it is interpreted as 'b' or V, depending on whether it is for a printer or a punch; no error indication is returned.

# INDEX

Indexes for reference manuals are consolidated in *OS/VS1 Master Index*, GC24-5104.

For additional information about any subject listed below, refer to the publications that are listed under the same subject in the master index.

## A

concatenation
  defined 145-146
  of data sets on RPS devices 146
  of partitioned data sets 145-146
  of sequential data sets 145-146
  of unlike data sets 145
  restriction with partitioned data sets 104
control buffer (see forms control buffer)
control character
  (see also CNTRL, PRTOV)
  ANSI 19,23,83,91,160
  carriage 18,26,159-160
  code 159-160
  explained 26
  format-D 23
  format-F 18,19
  format-U 25
  format-V 21
  machine 83,91,159-160
  specifying 82,159-160
  with fixed-length records 18,19
  with undefined-length records 25
  with variable-length records 21
control section, dummy (DSECT) 49
control volume 147
count area 28
  count-data format 28
  count-key-data format 28
  in device overhead 134
  in ISAM index entry format 108
CREATEV command, use of 133
creating a new generation 150-152
cross reference table with direct data sets 126
CSECT statement, use of with DCBD macro 49
cylinder
  allocation by 134
  capacity 16,135
  index
    calculating space requirements for 136-139
    definition 107-108
  overflow
    calculating space for 137,139
    defined 106,109
    specifying size
      via CYLOFL parameter 137
CYLOFL (cylinder overflow) operand of DCB macro
      when allocating ISAM data set 137
      when creating ISAM data set 118

# D

D-format records (see format-D records)
data access techniques (see access techniques)
data chaining 77
data check
  effect on BLKSIZE 18,83
  with SETPRT macro 87
data control block (DCB)
  ABEND exit
    description 44-47
    when available 34
    where specified 34
  attributes of, determining 31-33
  changing an address in 49,50
  creation by DCB macro instruction 15,30
  description 30-31
  dummy control section 49

exit
  description 33,42
  when available 34
  when used by SYSIN/SYSOUT 144
  where specified 34
fields 31-33
initial setting of 50
modifying 30,49-50
operand of DD statement 32
primary sources of information 30-31
sequence of completion 31
use 15
when sharing a data set 50
data definition name (ddname) field of DD statement 32
data definition (DD) statement
  fields 32-33
  relationship to DCB 30-31
  relationship to JFCB 30-31
  use 15
data errors 35-37
  (see also SYNAD routine)
data event control block (DECB)
  description of 61
  use of 94
data management, introduction to 13-53
data mode processing
  relationship with buffers 73
data overrun 77,78
data set
  characteristics 13-26
  description 31-33
  disposition (DISP) operand
    cataloging 143
    description 33
    overridden by OPEN macro 68,69
  identification 15
  label (DSCB) 15,16,155-158
    (see also magnetic tape volumes; data set control
      block; labels, direct-access)
  label (LABEL) field of DD statement 33
  like characteristics 145
  name 15
  name (DSNAME) field 32
  organization 14
    (see also BDAM, BISAM, BPAM, BSAM, QISAM,
      and QSAM data sets)
  organiztion (DSORG) operand of DCB macro 32
  record formats 17-26
  routing through the input/output stream 143-145
  security 13,153
  sharing 50-53
  space allocation on direct-access volumes 133-143
    estimation 134-136
    for a direct data set 126
    for indexed sequential data sets 136-138
    for MSS volume 134
    for partitioned data sets 136
    specifying 133-134
  storage 15,16
    direct-access 16
    magnetic-tape 16,17
  SYSIN 143-145
  SYSOUT 143-145
  unlike characteristics 145

# K

key
    class  116
    for direct-access devices  28
    for indexed sequential data sets  106-108
    protection  106
    relative key position (RKP) for indexed sequential data
        set  112,113,118
    use of when adding records to indexed sequential data
        set  109-111
    use of when maintaining an indexed sequential data
        set  111
    use of when retrieving records from an indexed sequential
        data set  120-125
KEYLEN operand of DCB macro
    description  32
    for direct-access device  85
    for direct data set  128
KN see WRITE KN)
KU (see READ KU)

# L

label exits  38-41
labels, data set  15-16,31,33
    (see also magnetic-tape volumes; labels, direct-access)
labels, direct-access
    data set control block  155-158
    format  155-158
    user label groups  157-158
    volume label group  155-157
LABEL parameter of DD statement
    description  33
    specifying password protection  153
    specifying standard labels  39
LEAVE option
    for close processing  67
    for concatendated data sets  145
    for end-of-volume processing  68,69
    for forced end-of-volume processing  69
length checking  18
link field  109,113-115
load mode for QISAM
    in SYNAD routine  37
    when sharing a DCB  52
load mode for BDAM when sharing data sets  52
loading an indexed sequential data set  118
locate mode processing
    defined for buffering  73-74
    example with exchange buffering  78-79
    example with simple buffering  75-77
    relationship with buffers  73
    to process records that exceed 32,760 bytes  23
    to update a member with QSAM  105
    with GET macro instruction
        creating a sequential data set, coding example  92-94
        exchange buffering  78,79
        simple buffering  74-76
    with parallel input processing  56
    with PUT macro instruction
        creating a sequential data set, coding example  92
        simple buffering  74-76,93-94
look-ahead buffering  55
LPALIB, restriction for checkpoint  43

LRECL operand of DCB macro
    described  32
    device dependence  90
    restriction when calculating record length  96
    to process records that exceed 32,760 bytes  23
    with BDAM  128
    with BSAM  96
    with ISAM
        buffer requirements  115
        data set creation  118
    with PUT macro  56
    with SYSOUT data set  144

# M

machine code control character  82,91,159,160
MACRF (macro instruction form) operand of DCB macro
    described  33
    device independence  89
    dynamic buffering  122
    for BDAM  128
    processing mode  73
    relationship with WAIT macro  60,61
    to update a member using QSAM  105
    when sharing a data set  51,52
magnetic-tape volumes
    defined  16,17
    density  83
    labels
        American National Standard  16,17
        none  16
        nonstandard  16
        standard  16
        user  38-39
    organization  16,17
    positioning  16
        during close processing  66-67
        during end-of-volume processing  68-69
    record format  17-26,83
    serial number  16,17
    tapemarks  17
master index  16,107
MBBCCHHR (see actual address)
modes, processing (see data mode; locate mode; move mode;
        substitute mode)
modifying the data control block  30,49,50
move mode processing
    defined for buffering  73
    relationship with buffers  73
    with exchange buffering  77
    with GET macro instruction
        creating a sequential data set  92
        simple buffering  74-76,92
    with parallel input processing  56
    with PUT macro instruction
        creating a sequential data set  92
        simple buffering  74-76,92
MSVGP parameter on JCL statemnt  134
MSWA operand of the DCB macro  115,116
multivolume data set  68
    processing for QISAM  71
multitasking mode, sharing data sets  52,64
multivolume data set
    with NOTE macro  87

# N

names
    data set  15
    generation data group  16,148,149
NCP (number of channel programs) operand of the DCB
       macro  58,91,144
nonstandard tape labels  16
note list  99
NOTE macro instruction
    description  87
    restriction with
      BSP macro  87
      multivolume data sets  87
      search direct operation  91
    updating a sequential data set  93
    use with partitioned data set
      updating  104
NTM operand  109
null segment  23

# O

offset reading  59
OMR (see optical mark read)
OPEN macro instruction
    considerations for  64
    description  65-66
    for parallel input processing  56-58
    for simultaneous opening of multiple data sets  63
    for updating a sequential data set  93
    functions  30,65,66
    restriction with search direct  91
    used for more than one data set  64
    volume positioning for EOV  68
opening a data set  63-66
OPTCD operand of the DCB macro
    device dependence  90
    with ASCII tapes (OPTCD=Q)  55
    with BDAM  127,128
    with ISAM  117
    request user totaling (OPTCD=T)  41
OPTCD=H (embedded checkpoints, DOS tapes)  86
OPTCD=M (master index)  109
OPTCD=T (user totaling)  41
OPTCD=Z (search direct option)  91
OUTIN option
    OPEN macro  65
    when opening data set  31
    when using POINT macro  88
OUTINX option
    when opening data set  31,65
output mode
    defined  74
OUTPUT option
    OPEN macro  65
    when opening data set  31
    when using POINT macro  88
output stream  143-145
overflow
    area  106,109
    chain  109

    cylinder (see cylinder overflow)
    independent area  109
    PROTV macro  86
    records  109
    track
      description  28
      effect on chained scheduling  90
      restriction on BSP macro instruction  87
      restriction with BDAM  131
      restriction with parallel input processing  56
      restriction with RPS feature  95
overlap of input/output
    performance improvement  90
    with basic access technique  58
    with partitioned data sets  105
    with sequential data sets  94
    with queued access technique  55

# P

paging environment, related control block group  58
paper-tape reader
    described  84
    effect on chained scheduling  90
    record format with  84
    with a SYNAD routine  36
parallel
    data access block (PDAB)  42,56,57
    input processing  42,56-58
parameter list
    contents of  46
    use of by DCB ABEND exit routine  45-47
partitioned data set (see BPAM data set)
PASSWORD data set  153
password protection  153
PC (card punch) record format  84
PCI flag  90
PDABD DSECT  57
PDAB (parallel data access block)  42,56,57
PDS (see BPAM data set)
performance improvement  90
POINT macro instruction
    description  88
    in EODAD routine  34
    restriction with
      BSP macro  87
      multivolume data sets  88
      search direct operation  91
    updating a partitioned data set  104
    updating a sequential data set  93
prefix, block (see block prefix)
prefix, key  116
prime data area
    description  106,107
    space allocation for  136,138-139
printer
    overflow (PRTOV macro)  86
    record format with  85
    restriction with chained scheduling  91
program, describing the processing  33-48
PROTECT macro instruction  153
PRTOV macro instruction
    description  86
    device dependent  89
    when macro will not function  86
PT (see paper-tape reader)

PUT macro instruction
    description 55-56
    exchange buffering 77-80
    locate mode 73-76
    used to create a sequential data set, coding
        example 92-94
    with format-U records 25
    with indexed sequential data set 109-111
    with simple buffering 74-76
    with spanned records 23
        (see also data mode processing; locate mode
            processing; move mode processing; substitite mode
            processing)
PUTX macro instruction
    description 56
    device independence 89
    for QISAM 122
    UPDAT mode 77
    updating a sequential data set 93,94
    when sharing a data set 51
    with exchange buffering 78,80,81
    with format-U records 25-26
    with simple buffering 74-77
    with spanned records 23
        (see also output mode; update mode)

## Q

QISAM data set
    (see also indexed sequential data set)
    EODAD routine 34
    scan mode 122
    sharing 51,52
    SYNAD routine 36,37
    using common buffer pool 71
QSAM
    (see also queued access technique)
    creating a BDAM data set 128
    parallel input exit 42
    parallel input processing 56-58
    performance improvement 90
    restriction with
        spanned records 21
    spanned variable-length records 21-23
    SYSIN/SYSOUT data sets 144
    to update a directory 101
    to update a member 105
    user labels 38
    user totaling 41-42
    when sharing a data set 51,52
    with card punch 84
    with printer 85
queued access technique
    buffer control 70,73-74
    defined 55
    introduced 14
    processing modes (see data mode processing; locate mode
        processing; move mode processing; substitute mode
        processing)

## R

RD (card reader) 84,85
RDBACK option 40
    opening magnetic tape volume 31
    restriction for variable-length records 65
    restriction with SYSIN/SYSOUT data sets 65
RDW (see record descriptor word)
read backward (SB operand of READ macro) 59
READ macro instruction
    description 59
    device independence 89
    in SYNAD routine 36
    restriction in EODAD routine 34
    supplying record length 96
    to enter EODAD routine 33
    to update existing records 121
    updating a partitioned data set 104
    updating a sequential data set 93,94
    when sharing a data set 51,52
    with basic access technique 58
    with format-U records 26
    with KU (key, update)
        in coding example 123
RECFM operand of DCB macro
    description 32
    for sequential data sets 82
    selecting 18,19
    with card punch 84
    with card reader 84
    with control character 82
    with direct-access storage device 85
    with magnetic tape 83
    with paper tape reader 84
    with printer 85
    with sequential organization 82
record blocking (see blocking)
record descriptor word (RDW) 21
    data mode exception for spanned records 21
    variable-length records format-D 23,24
    when replaced by segment descriptor word 22
record format 17-26
    fixed-length 17-20
    fixed-length for ASCII 17-20
    fixed-length standard 18-20
    spanned variable-length 21-23
    undefined-length 17
    variable-length 17,20-21
record length (LRECL) operand of the DCB macro 32
relative block address
    defined 29
    with direct data set 127,129
    with feedback option 127,128
relative generation name 148-150
relative key position operand of the DCB macro 112,113,118
relative track address
    defined 29
    with direct access 127,129
    with feedback option 127,128
RELEX macro instruction
    exclusive control 52,128
    when sharing data sets 52
RELSE macro instruction
    defined 80
    to terminate buffer processing 70
reorganization of indexed sequential data set 111

REREAD option 68,69
restart end-of-volume exit routine 43
restrictions
    on ASCII records
        block prefix 19,20,24
        on 7-track tape 83
    on chained scheduling with
        BDAM 90
        calculating record length 95,96
        CNTRL macro 90
        DOS checkpoint records 90
        exchange buffering 90
        format-D records 24
        paper-tape reader 90
        partitioned data set 104
        SKP option 35
        spooled data sets 90
        track overflow 90,95
        UPDAT operand 90
        updating a sequential data set 93,94
        2540 Card Read Punch 91
        3525 Card Punch 91
    on CNTRL macro
        with BSP macro 87
        with chained scheduling 90
        with DOS checkpoint records 86
    on data sets with same DDname 64
    on DCB usage 64-65
    on DCBD macro usage 49-50
    on DOS checkpoint records 86-88,90
    on format-D records with chained scheduling 24
    on high-level index in storage 116
    on NOTE macro with
        BSP macro 87
        multivolume data sets 87
        search direct operation 91
    on POINT macro with
        BSP macro 87
        multivolume data sets 88
        search direct operation 91
    on reading concatenated data sets backward 145
    on user label exit routines 38-41
    with search direct 91
resume load 110,111,118-120
retrieving a generation 152
return code
    with block count exit 44
    with user labels 40
RETURN macro
    relationship in SYNAD routine 35
REWIND option
    for CLOSE macro 66
    for FEOV macro 69
RKP (relative key position 112,113,118
RLSE parameter of DD statement 67
RORG1, RORG2, RORG3 fields of the DCB 111
routing data sets through the input/output stream 143-145
RPS (rotational position sensing) feature
    concatenating data sets on nonRPS devices 146
    restriction with track overflow records
        variable-length records 20
        when calculating record length 95
R0 record 28,128,133

# S

save area, user totaling 41-42
scan mode for QISAM
    in SYNAD routine 37
    issuing PUTX 122
scheduling of input/output streams 143
SDW (see segment descriptor word)
search direct for input 91
search option, extended 127
secondary storage (see data set storage; direct-access storage;
        magnetic-tape volumes)
security, data set 13,153
segment
    buffer 70,71
    control code 22
    descriptor word
        for spanned records 22,23
        indicating a null segment 23
    null 23
selecting an access method 62
SEP (separation, channel) 32
sequential data set
    (see also BPAM, BSAM, and QSAM data sets)
    creation 92
    concatenation 145-146
    extending 94
    retrieving 93
    updating 93-94
sequential organization
    defined 14
    device control 86-88
    device independence 88-89
SETL macro instruction 116-117
    in EODAD routine 34
    when sharing a data set 52
SETPRT macro instruction 86-87
SETPRT routine 44
sharing data sets 50-53
sharing DASDs 53
simple buffering
    description 74-77
    when defaulted to 78
    with parallel input processing 56,57
SKP error option 35
SMSI operand of the DCB macro 116
SMSW operand of the DCB macro 115,116
Sort/merge program
    record restriction 18
space allocation
    estimating requirements 134-136
    for a direct data set 126
    for an indexed sequential data set 136-138
    for an MSS volume 134
    for a partitioned data set 136
    specifying 133,134
SPACE parameter 32
spanned records
    basic direct access method 23
    considerations for 21-23
    restriction with parallel input processing 56
    restriction with search direct 91
    restriction with SYSIN data sets 23
    sequential access method 21
    variable-length 21

# 123

GC26-3874-0

# GLOSSARY OF ACRONYMS AND ABBREVIATIONS

The following terms are defined as they are used in this book. If you do not find the term you are looking for, refer to the index or to the *IBM Data Processing Glossary*, GC20-1699.

| | |
|---|---|
| **A** | ANSI control code (value of RECFM) |
| **ABE** | abnormal end (value of EROPT) |
| **ABEND** | abnormal end (macro instruction) |
| **ACC** | accept erroneous block (value of EROPT) |
| **AFF** | affinity (channel separation parameter of DD statement or unit affinity value of UNIT) |
| **ANSI** | American National Standards Institute |
| **ASCII** | American National Standard Code for Information Interchange |
| **ABSTR** | absolute track (value of SPACE) |
| **AL** | American National Standard Labels |
| **AUL** | American National Standard User labels (value of LABEL) |
| | |
| **B** | blocked records (value of RECFM) |
| **BCDIC** | binary coded decimal interchange code |
| **BDAM** | basic direct access method |
| **BDW** | block descriptor word |
| **BFALN** | buffer alignment (operand of DCB) |
| **BFTEK** | buffer technique (operand of DCB) |
| **BISAM** | basic indexed sequential access method |
| **BLDL** | build list (macro instruction) |
| **BLKSIZE** | blocksize (operand of DCB) |
| **BPAM** | basic partitioned access method |
| **BPI** | bits per inch |
| **BSAM** | basic sequential access method |
| **BSM** | backspace past tapemark and forward space over tapemark (operand of CNTRL) |
| **BSP** | backspace one block (macro instruction) |
| **BSR** | backspace over a specified number of blocks (records) (operand of CNTRL) |
| **BUFCB** | buffer pool control block (operand of DCB) |
| **BUFL** | buffer length (operand of DCB) |
| **BUFNO** | buffer number (operand of DCB) |
| **BUFOFF** | buffer offset (length of ASCII block prefix by which the buffer is offset; operand of DCB) |
| | |
| **CCW** | channel command word |
| **CONTIG** | contiguous space allocation (value of SPACE) |
| **CNTRL** | control (macro instruction) |
| **CPU** | central processing unit |
| **CSW** | channel status word |
| **CYLOFL** | number of tracks for cylinder overflow records (operand of DCB) |
| | |
| **D** | format-D (ASCII variable-length) records (value of RECFM) |
| **DA** | direct-access (value of DEVD or DSORG) |
| **DAU** | direct-access unmovable data set (value of DSORG) |
| **DCB** | data control block (control block name or macro instruction) |
| **DCBD** | data control block dummy section macro instruction |
| **DD** | data definition |
| **DEB** | data extent block |

| | |
|---|---|
| **DECB** | data event control block |
| **DEN** | magnetic tape density (operand of DCB) |
| **DEVD** | device-dependent (operand of DCB) |
| **DISP** | data set disposition (parameter of DD statement) |
| **DSCB** | data set control block |
| **DSORG** | data set organization (operand of DCB) |
| **EBCDIC** | extended binary coded decimal interchange code |
| **EODAD** | end-of-data set exit routine address (operand of DCB) |
| **EOF** | end-of-file |
| **EOV** | end-of-volume |
| **EROPT** | error options (operand of DCB) |
| **ESETL** | end sequential retrieval (QISAM macro instruction) |
| **EXCP** | execute channel program (macro instruction) |
| **EXLST** | exit list (operand of DCB) |
| **F** | fixed-length records (value of RECFM) |
| **FB** | fixed-length, blocked records (value of RECFM) |
| **FBS** | fixed-length, blocked, standard records (value of RECFM) |
| **FBT** | fixed-length, blocked records with track overflow option (value of RECFM) |
| **FCB** | forms control buffer |
| **FEOV** | force end-of-volume (macro instruction) |
| **FS** | fixed-length, standard records (value of RECFM) |
| **FSM** | forward space past tapemark and backspace over tapemark (operand of CNTRL) |
| **FSR** | forward space over a specified number of blocks (records) (operand of CNTRL) |
| **GCR** | group coded recording |
| **GL** | GET macro, locate mode (value of MACRF) |
| **GM** | GET macro, move mode (value of MACRF) |
| **HA** | home address |
| **I/O** | input/output |
| **INOUT** | input then output (operand of OPEN) |
| **IOB** | input/output block |
| **IPL** | initial program load |
| **IRG** | interrecord gap |
| **IS** | indexed sequential (value of DSORG) |
| **ISAM** | indexed sequential access method |
| **ISU** | indexed sequential unmovable (value of DSORG) |
| **JCL** | job control language |
| **JFCB** | job file control block |
| **JFCBE** | job file control block extension for 3800 printer |
| **KEYLEN** | key length (operand of DCB) |
| **LPA** | link pack area |
| **LPALIB** | link pack area library |
| **LRECL** | logical record length (operand of DCB) |
| **M** | machine control code (value of RECFM) |
| **MACRF** | macro instruction form (operand of DCB) |
| **MOD** | modify data set (value of DISP) |
| **MSHI** | main storage for highest-level index (operand of DCB) |
| **MSS** | Mass Storage System |

| | |
|---|---|
| MSVC | Mass Storage Volume Control |
| MSWA | main storage for work area (operand of DCB) |
| | |
| NCP | number of channel programs (operand of DCB) |
| NOPWREAD | no password to read a data set (value of LABEL) |
| NRZI | non-return-to-zero-inverted (tape recording mode) |
| NSL | nonstandard label (value of LABEL) |
| NTM | number of tracks in cylinder index for each entry in lowest level of master index (operand of DCB) |
| | |
| OMR | optical mark read |
| OPTCD | optional services code (operand of DCB) |
| OS/VS | operating system/virtual storage |
| OUTIN | output then input (operand of OPEN) |
| | |
| PCI | program-controlled interruption |
| PDAB | parallel data access block |
| PDS | partitioned data set |
| PE | phase encoding (tape recording mode) |
| PL | PUT macro, locate mode (value of MACRF) |
| PM | PUT macro, move mode (value of MACRF) |
| PO | partitioned organization (value of DSORG) |
| POU | partitioned organization unmovable (value of DSORG) |
| PRECL | physical record length (field of DCB) |
| PRTSP | printer line spacing (operand of DCB) |
| PS | physical sequential (value of DSORG) |
| PSU | physical sequential unmovable (value of DSORG) |
| | |
| QISAM | queued indexed sequential access methods |
| QSAM | queued sequential access method |
| | |
| RCE | read column eliminate |
| RDBACK | read backward (operand of OPEN) |
| RDW | record descriptor word |
| RECFM | record format (operand of DCB) |
| RKP | relative key position (operand of DCB) |
| RLSE | release unused space (DD statement) |
| RPS | rotational position sensing |
| | |
| S | standard format records (value of RECFM) |
| SDW | segment descriptor word |
| SEP | separation (channel separation parameter of DD statement or unit separation value of UNIT) |
| SER | volume serial number (value of VOLUME) |
| SETL | set lower limit of sequential retrieval (QISAM macro instruction) |
| SF | sequential forward (operand of READ or WRITE) |
| SK | skip to a printer channel (operand of CNTRL) |
| SKP | skip erroneous block (value of EROPT) |
| SL | IBM standard labels (value of LABEL) |
| SMSI | size of main-storage area for highest-level index (operand of DCB) |
| SMSW | size of main-storage work area (operand of DCB) |
| SP | space lines on a printer (operand of CNTRL) |
| SS | select stacker on card reader (operand of CNTRL) |
| SUL | IBM standard and user labels (value of LABEL) |
| SVC | supervisor call |
| SVCLIB | supervisor call library |
| SYNAD | synchronous error routine address (operand of DCB) |

| | |
|---|---|
| **SYSIN** | system input stream |
| **SYSOUT** | system output stream |
| **T** | track overflow option (value of RECFM) |
| **TIOT** | task I/O table |
| **TRTCH** | track recording technique (operand of DCB) |
| **U** | undefined length records (value of RECFM) |
| **UCS** | universal character set |
| **UHL** | user header label |
| **UTL** | user trailer label |
| **V** | format-V (variable-length) records (value of RECFM) |
| **VB** | variable-length, blocked records (value of RECFM) |
| **VBS** | variable-length, blocked, spanned records (value of RECFM) |
| **VS** | virtual storage or variable-length, spanned records |
| **VTOC** | volume table of contents |

OS/VS1 Data Management Services Guide
GC26-3874-0

Your comments about this publication will help us to improve it for you.
Comment in the space below, giving specific page and paragraph references
whenever possible. All comments become the property of IBM.

Please do not use this form to ask technical questions about IBM systems and
programs or to request copies of publications. Rather, direct such questions or
requests to your local IBM representative.

If you would like a reply, please provide your name and
address (including ZIP code).

**Fold on two lines, staple, and mail.** No postage necessary if mailed in the U.S.A. (Elsewhere,
any IBM representative will be happy to forward your comments.) Thank you for your
cooperation.

GC26-3874-0

First Class Permit
Number 6090
San Jose, California

## Business Reply Mail

No postage necessary if mailed in the U.S.A.

Postage will be paid by:

**IBM Corporation**
**P.O. Box 50020**
**Programming Publishing**
**San Jose, California 95150**

IBM
®

OS/VS1 Data Management Services Guide (File No. S370-30) Printed in U.S.A. GC26-3874-0