IBM

Tape and Diskette Device Programming

Version 4

AS/400e

Tape and Diskette Device Programming

Version 4

Second Edition (May 1999)

Contents

Figures

About Tape and Diskette Device Programming (SC41-5716-01)

This book describes the tape and diskette media supported by the AS/400 system and the Operating System/400 licensed program. Characteristics and programming use of tape and diskette files are described.

Data management concepts and programming considerations for input and output spooling are described in the *Data Management* book. Printer device file characteristics and programming use is described in the *Printer Device Programming* book. The book describes display file device characteristics and programming use in the *Application Display Programming* book. The book describes information about tasks that are performed with an automated tape library (ATL) in the *Automated Tape Library Planning and Management* book. The tape management exit program provides the function to monitor and control the use of volumes and devices that are used by the operating system. The book describes the exit program in the *System API Reference* book.

For more information about other AS/400 publications, see the following:

- The *AS/400 Information Directory*, a unique, multimedia interface to a searchable database containing descriptions of titles available from IBM or from selected other publishers.

For a list of publications related to this guide, see the "" on page 85.

Who Should Read This Book

This book is intended primarily for the application programmer.

Before using this guide, you should be familiar with general programming concepts and terminology, and have a general understanding of the AS/400 system and OS/400 program.

AS/400 Operations Navigator

AS/400 Operations Navigator is a powerful graphical interface for Windows clients. With AS/400 Operations Navigator, you can manage and administer your AS/400 systems from your Windows desktop.

You can use Operations Navigator to manage communications, printing, database, security, and other system operations. Operations Navigator includes Management Central for managing multiple AS/400 systems centrally.

Figure 1 on page viii shows an example of the Operations Navigator display:
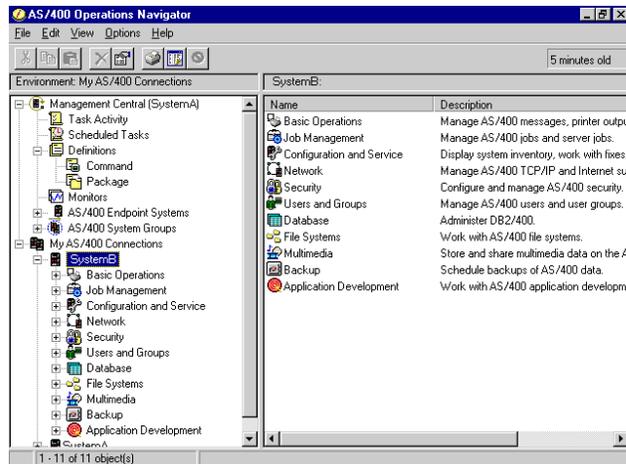
Figure 1. AS/400 Operations Navigator Display

This new interface has been designed to make you more productive and is the only user interface to new, advanced features of OS/400. Therefore, IBM recommends that you use AS/400 Operations Navigator, which has online help to guide you. While this interface is being developed, you may still need to use a traditional emulator such as PC5250 to do some of your tasks.

Installing Operations Navigator

To use AS/400 Operations Navigator, you must have Client Access installed on your Windows PC. For help in connecting your Windows PC to your AS/400 system, consult *Client Access Express for Windows - Setup*, SC41-5507-00.

AS/400 Operations Navigator is a separately installable component of Client Access that contains many subcomponents. If you are installing for the first time and you use the **Typical** installation option, the following options are installed by default:

- Operations Navigator base support
- Basic operations (messages, printer output, and printers)

To select the subcomponents that you want to install, select the **Custom** installation option. (After Operations Navigator has been installed, you can add subcomponents by using Client Access Selective Setup.)

1. Display the list of currently installed subcomponents in the **Component Selection** window of **Custom** installation or Selective Setup.
2. Select AS/400 Operations Navigator.
3. Select any additional subcomponents that you want to install and continue with **Custom** installation or Selective Setup.

After you install Client Access, double-click the **AS400 Operations Navigator** icon on your desktop to access Operations Navigator and create an AS/400 connection.

Prerequisite and related information

Use the AS/400 Information Center as your starting point for looking up AS/400 technical information. You can access the Information Center from the AS/400e Information Center CD-ROM (English version: *SK3T-2027*) or from one of these Web sites:

```
http://www.as400.ibm.com/infocenter
http://publib.boulder.ibm.com/pubs/html/as400/infocenter.htm
```

The AS/400 Information Center contains important topics such as logical partitioning, clustering, Java, TCP/IP, Web serving, and secured networks. It also contains Internet links to Web sites such as the AS/400 Online Library and the AS/400 Technical Studio. Included in the Information Center is a link that describes at a high level the differences in information between the Information Center and the Online Library.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other AS/400 documentation, fill out the readers' comment form at the back of this book.

- If you prefer to send comments by mail, use the readers' comment form with the address that is printed on the back. If you are mailing a readers' comment form from a country other than the United States, you can give the form to the local IBM branch office or IBM representative for postage-paid mailing.
- If you prefer to send comments by FAX, use either of the following numbers:
  - United States and Canada: 1-800-937-3430
  - Other countries: 1-507-253-5192
- If you prefer to send comments electronically, use one of these e-mail addresses:
  - Comments on books:
    - RCHCLERK@us.ibm.com
    - IBMMAIL, to IBMMAIL(USIB56RZ)
  - Comments on the AS/400 Information Center:
    - RCHINFOC@us.ibm.com

Be sure to include the following:
- The name of the book.
- The publication number of the book.
- The page number or topic to which your comment applies.

Chapter 1. Introduction

Device files are files that provide access to externally attached devices such as tapes, diskettes, printers, displays, spools, and other systems that are attached by a communications line. The two device files that are described in this manual are:

- **Tape files** which allow access to data files on tape devices
- **Diskette files** which provide access to data files on the diskette unit

Each file type has its own set of unique characteristics that determines the use and capabilities the file can provide. This manual describes the characteristics and use of tape and diskette device files for application programs. For more information about display files, printer files, and spooled files, refer to the *Application Display Programming*, *Printer Device Programming* ,and *Data Management* books, respectively.

When programs use a device file, a name represents the device file. That name identifies both the file description and, for some file types, the data itself. This manual helps you understand the following aspects of tape and diskette files so you can use their full capabilities:

- Usage characteristics
- Configuration descriptions
- Error handling methods
- Usage in high-level language programs

Chapter 2. Tape Support

The AS/400 system supports the following tape media for creating backup copies and offline storage of information, or for transferring information to other devices or systems:

- 1/2-inch reel
- 1/2-inch cartridge
- 8-mm cartridge
- 1/4-inch cartridge
- 1/4-inch mini-cartridge

Tape is particularly useful for storing large amounts of data.

For information about how to use tape devices for save and restore operations, see the *Backup and Recovery*, SC41-5304-03 book.

Related CL Commands

The following commands are available to help maintain and use tapes. The *CL Reference (Abridged)*, SC41-5722-03 manual contains detailed descriptions of these commands.

Tape Configuration Description Commands

CHGCTLTAP
> Change Controller Description (Tape): The command changes the controller description for a tape controller.

CHGDEVTAP
> Change Device Description (Tape): The command changes the device description for a tape device.

CHGDEVMLB
> Change Device Media Library (Tape): The command changes the device description for a media library device.

CFGDEVMLB
> Configure Device Media Library (Tape): The command configures the connection between the media library device and the communication interfaces of the robotics.

CRTCTLTAP
> Create Controller Description (Tape): The command creates a controller description for a tape controller.

CRTDEVTAP
> Create Device Description (Tape): The command creates a device description for a tape device.

CRTDEVMLB
> Create Device Media Library (Tape): The command creates a device description for a media library device.

DLTCTLD
> Delete Controller Description: The command deletes a controller description.

DLTDEVD

Delete Device Description: The command deletes a device description.

DSPCTLD

Display Controller Description: The command displays a controller description.

DSPDEVD

Display Device Description: The command displays a device description.

DSPLANMLB

Display LAN Media Library (Tape): The command displays the LAN information necessary to configure a library manager.

Tape Device File Commands

CHGTAPF

Change Tape File: The command changes certain attributes of a tape device file.

CRTTAPF

Create Tape File: The command creates a tape device file that is used to read and write records on tape.

DLTF   Delete File: The command deletes files.

DSPFD

Display File Description: The command displays the current characteristics of a file.

OVRTAPF

Override with Tape File: The command temporarily changes a tape file or tape file attributes that are specified in a program.

Other Tape Support Commands

ADDTAPCTG

Add Tape Cartridge: The command adds the specified cartridge identifiers to a usable category.

CHGJOBMLBA

Change Job MLB Attributes: The command allows a user to change the media library resource allocation attributes for a job.

CHGTAPCTG

Change Tape Cartridge: The command changes the specified cartridge from any category to the specified category.

CHKTAP

Check Tape: The command searches a tape volume for a specific volume identifier or file label.

CPYFRMTAP

Copy from Tape: The command copies records from a tape file to an output file or a printer.

CPYTOTAP

Copy to Tape: The command copies records to a tape file from a physical, logical, tape, diskette, or spooled inline data file.

CRTTAPCGY

Create Tape Category: The command creates a user-defined category name and assigns it to a system name.

DLTTAPCGY

Delete Tape Category: The command deletes a user-defined category name that was previously created with the Create Tape Category (CRTTAPCYG) command. The command will not delete the category if a cartridge currently uses that category.

DMPTAP

Dump Tape: The command dumps label information, data blocks, or both, from a tape with or without a label.

DSPTAP

Display Tape: The command displays volume label information,file label information, and saved object information for standard label tapes. The command displays both the volume type and density for volumes without labels.

DSPTAPCGY

Display Tape Category: The command allows the user to display the categories that are defined through the Create Tape Category (CRTTAPCGY) command.

DSPTAPCTG

Display Tape Cartridge: The command displays the attributes of tape cartridges.

DSPTAPSTS

Display Tape Status: The command displays the slot information for the library device and tape device information for the tape devices that are attached to the library device.

DUPTAP

Duplicate Tape: The command copies the contents of one tape to another.

INZTAP

Initialize Tape: The command initializes tapes, with or without labels, or clears all the data on the tape from the load point to the end-of-tape marker.

RMVTAPCTG

Remove Tape Cartridge: The command removes the specified cartridge identifiers from their current category, or the specified category, and places them in the eject (*EJECT) category. The command can move the specified cartridge to the Convenience I/O station or to the High Capacity Output station. The eject category is not a valid category for I/O operations. The Eject category cartridges are disallowed in the tape device.

SETTAPCGY (Set Tape Category)

Set Tape Category: The command allows the user to set a category on a tape device within the specified media library.

WRKMLBRSCQ

The command allows a user to work with the resource allocation requests for the specified media library device.

WRKTAPCTG

Work with Tape Cartridges: The command allows the user to work with a list of tape cartridges.

Initializing Tapes

Before use you must initialize all tapes. Use the Initialize Tape (INZTAP) command to initialize tapes, with or without labels. Sometimes you use the Initialize Tape (INZTAP) command to clear all data on the tape from the load point to the end-of-tape marker.

Note: If the tape device is an 8-mm cartridge device and specifies the command, CLEAR(*YES), then the operation will take over 3.5 hours. The program erases the tape from the beginning-of-tape to the end-of-tape.

The following example initializes the tape volume that is loaded on device TAP01 to standard label format.

```
INZTAP  DEV(TAP01) NEWVOL(BACKUP)
DENSITY(*FMT3490E)
```

Initialize the tape volume with the volume ID BACKUP, you specify the character code EBCDIC(by default) and set the tape format to *FMT3490E.

To convert from one language to another as in the examples below:
- EBCDIC to ASCII
- ASCII to EDCDIC

The AS/400 system uses data conversion tables that are derived from the American National Standards Institute, Inc, Document ANSI X3.26-1970.

Tape Labeling

The following series of diagrams provides a basic description of standard tape labeling used for the AS/400 system.

In Figure 2, the INZTAP command gives the tape a volume label (marked VOL1) and writes two tape marks (TM).
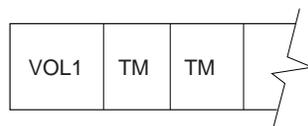
RSLH146-0

Figure 2. Volume Label and Tape Marks

When a high-level language program opens a tape file the system does the following procedures:

1. Writes over the two tape marks that follow the VOL label with header labels HDR1 and HDR2.
2. It adds a single tape mark that follows the header labels.

Each header label is 80 bytes-long. The first header label contains such information as the file name and date. The second header label specifies information such as record and block lengths, record block format, and buffer offset (for ASCII files).

When a high-level language program writes data to tape the system writes the data to the tape after the tape mark. Reaching the end of the file the system writes a

tape mark and two end-of-file labels on the tape. The end-of-file labels contain the same information as the header labels except that the first end-of-file label (EOF1) includes the block count for the file.

Two tape marks follow the end-of-file labels as shown in the diagram below.

VOL1	HDR1	HDR2	TM	Data

TM	EOF1	EOF2	TM	TM

RSLH147-0
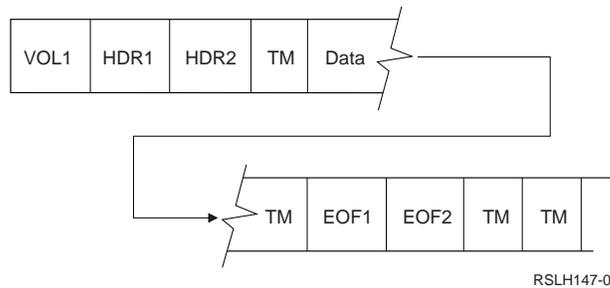
When the high-level language program adds a second file to the tape, the system creates a header label (HDR1) for the new file. This header label (HDR1) for the new file writes over the second tape mark following the end-of-file labels. The second header label, another tape mark, and the file data follow the new header label (HDR1) as illustrated below.

TM	EOF1	EOF2	TM	HDR1	HDR2	TM	Data

RSLH148-0
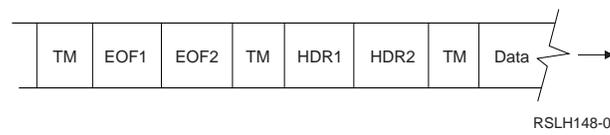
When the tape drive reaches the end of the physical tape, the system writes two tape marks that follow the end-of-volume labels. If the file is not complete, it continues on a second volume, that specifies the tape as volume 2 of the file.

TM	EOV1	EOV2	TM	TM

RV3W126-0

Tape Data Files

Store data files on a tape in the following ways:
- Single volume tape data file: A file that is contained on one volume of tape
- Multivolume-tape data files: Files that are contained on more than one volume of tape.
- Multifile volumes: Volumes of tape that contain more than one data file

If you use multivolume-tape data files, then follow these conventions:
- The labels on each volume must be consistent. You cannot have standard labeled tapes and unlabeled tapes in the same tape group.
- Write all volumes and density in the same character code (EBCDIC or ASCII).
- Each tape in the group must have the same record format, block length, and record length.

- If you specify more than one tape device, then place the volumes on the devices in the sequence that is specified in the tape device file. Refer to the following reference: Figure 3.For example:
  - The data file consists of four volumes such as VOL01, VOL02, VOL03, and VOL04.
  - The tape devices specified, in order, are TAPE01, TAPE02, and TAPE03.

  Then place the volumes on a tape device as follows: VOL01 on TAPE01, VOL02 on TAPE02, VOL03 on TAPE03, and VOL04 on TAPE01.

```
┌─────────┐   ┌─────────┐   ┌─────────┐
│         │   │         │   │         │
│ TAPE01  │   │ TAPE02  │   │ TAPE03  │
│         │   │         │   │         │
└─────────┘   └─────────┘   └─────────┘
  VOL01────────►VOL02────────►VOL03───┐
┌►VOL04                               │
└─────────────────────────────────────┘
```
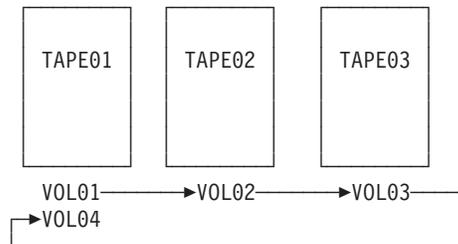
Figure 3. Multivolume-Tape-Data-File Sequence Using Three Tape Devices

If you use volumes, in reverse order, by reading backwards, then VOL04 is on TAPE01, VOL03 on TAPE02, VOL02 on TAPE03, and VOL01 on TAPE01.

```
┌─────────┐   ┌─────────┐   ┌─────────┐
│         │   │         │   │         │
│ TAPE01  │   │ TAPE02  │   │ TAPE03  │
│         │   │         │   │         │
└─────────┘   └─────────┘   └─────────┘
  VOL04────────►VOL03────────►VOL02───┐
┌►VOL01                               │
└─────────────────────────────────────┘
```
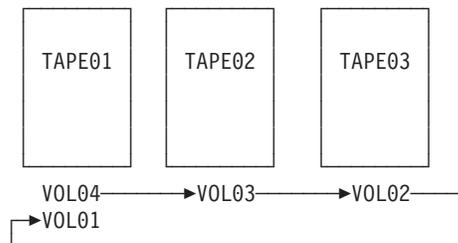
Tape Configuration Descriptions and Device Files

To access data from a tape device on the system, the following objects must exist:

- First, a device descriptionmust exist for each tape device to describe the device to the system.You specify the tape device description using the Create Device Description (Tape) (CRTDEVTAP) command. The tape device description contains information such as device address, device name, device type, model number, and features. For some tape devices, a tape controller description must also exist.
- Second, a device file must exist for thetape device. Create the tape device files by using the Create Tape File (CRTTAPF) command. Tape device files describe how a device presents input data to a program, or how a program presents output data to a device. Do not confuse tape device files with the actual data files on the tape volumes. For processing volumes which contain data files, the tape device files provide a link between the application program and the tape device.

It is not necessary to have a separate device file for each tape device. Your application program can use a single device file for several different tape devices

through the use of an Override Tape File (OVRTAPF) command. You can associate
any number of device files with one device.

Note: You must vary on the configuration descriptions before using them. See the
Local Device Configuration, SC41-5121-00 book for information about
varying on configuration descriptions.

IBM-Supplied Tape Device Files

You can use the following tape device files that are shipped with the operating
system:

- QTAPE (tape file)
- QTAPSRC (tape source file)

These files are all program-described as data files in library QGPL. The record
format names are the same as the file names. The files contain default values for
most parameters.

You can create additional tape device files to fit your needs. For example, you can
create an additional tape device file to contain the specific volume and label
information for a tape data file that several programs can use.

Example of Creating a Tape Device File

In the following example, the program creates a tape device file, TAP05 in library
QGPL, for output that is written to tape:

```
CRTTAPF  FILE(QGPL/TAP05) DEV(TAP01)
REELS(*SL) SEQNBR(3)
CODE(*EBCDIC) ENDOPT(*UNLOAD)
```

The program specifies the tape REELS parameter with the value *SL, indicating
that the tape uses standard labels. The device name is TAP01. The program writes
a file at sequence number three on the tape (SEQNBR parameter), in EBCDIC
code (CODE parameter), and unloads it after it has processed (ENDOPT
parameter).

Extending Files on Tape

Tape data files on 1/2-inch tapes can be extended using the EXTEND parameter on
the CRTTAPF, CHGTAPF, and OVRTAPF commands. The system does not support
extending data files for 1/4-inch or 8-mm cartridge tape devices.

When you extend a file, any existing tape data following the specified file on the
tape is no longer accessible by the system.

In the following example, a tape contains four files: FILE1, FILE2, FILE3, and
FILE4. If FILE2 is extended, FILE3 and FILE4 are no longer accessible.

Note: If you specify EXTEND(*YES *CHECK) on the OVRTAPF command, the
expiration date of the file (FILE3) that is following the extension of (FILE2) is
checked. The expiration date will be checked *before* extending the file
(FILE2). However, the expiration dates of any remaining files (FILE4) are not
checked, even if EXTEND(*YES *CHECK) is specified.

```
┌──────┬──────┬──────┬──────┬────
│      │      │      │      │
│FILE1 │FILE2 │FILE3 │FILE4 │
│      │      │      │      │
└──────┴──────┴──────┴──────┴────
```

```
┌──────┬──────────────┬──────────
│      │              │  .    .
│FILE1 │ FILE2──────► │  .    .
│      │              │  .    .
└──────┴──────────────┴──────────
```

Specifying Tape Device File Parameters

The description of the tape device file record is in the application program that uses the tape information. The system views each record as one field with a length equal to the record length.

The following section lists considerations for parameters that are specified on the CRTTAPF, CHGTAPF, and OVRTAPF commands.

DEV   The name of the device description for a tape device file that identifies the devices the file can access.

VOL   The volume identifiers of the tapes that are used for the device file may be specified using the VOL parameter on the CRTTAPF, CHGTAPF, and OVRTAPF commands. The volume identifiers may contain from 1 to 6 alphanumeric characters.

REELS
The REELS parameter specifies both the number of tapes that will contain the data file, and the type of label processing that is used by those tapes. Ignore the reel numberduring output processing or specifying a volume list. Ignore the reel number if you specify standard label processing (by using *SL on the REELS parameter).

If some of the file labels are incorrect, specify bypass label processing (*BLP). The system will check each reel for a volume label that begins with the characters VOL1. The system will ignore most other volume label information and the file labels on the tape.

For bypass label processing, each data file on the tape must contain a header label and either an end-of-file trailer label or anend-of-volume trailer label.

SEQNBR
The SEQNBR parameter specifies the sequence number of the data file on tape. The data files are numbered consecutively across all the volumes they occupy, starting with sequence number 1 for the first data file on the first volume. (Valid values for the sequence number range from 1 to16 777 215.) Figure 4 on page 11 shows how to number files for labeled volumes that contain more than one file and contain multivolume tapes (FILEB on three volumes):
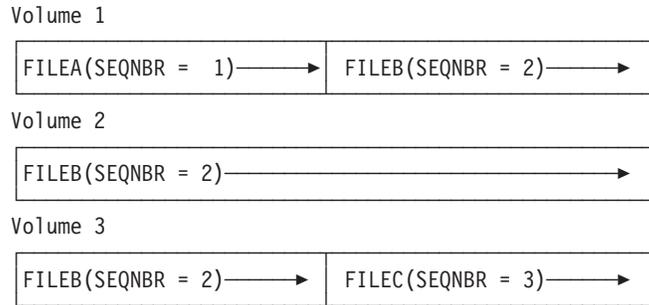
```
Volume 1
┌────────────────────────────┬─────────────────────────────┐
│ FILEA(SEQNBR =  1)─────────►│ FILEB(SEQNBR = 2)─────────►  │
└────────────────────────────┴─────────────────────────────┘
Volume 2
┌──────────────────────────────────────────────────────────┐
│ FILEB(SEQNBR = 2)───────────────────────────────────────► │
└──────────────────────────────────────────────────────────┘
Volume 3
┌────────────────────────────┬─────────────────────────────┐
│ FILEB(SEQNBR = 2)─────────► │ FILEC(SEQNBR = 3)─────────►  │
└────────────────────────────┴─────────────────────────────┘
```

Figure 4. Data File Sequence Number on Multivolume Tapes

The sequence number specified for new standard labeled data files on tape must correspond to the physical sequence number of data files on the tape. (Specify the sequence number by the SEQNBR parameter on the CRTTAPF, CHGTAPF, and OVRTAPF commands.) This means that if files 1 and 2 exist on the tape, the next data file created must have a sequence number of 3. To create a new data file on a tape that contains the last volume of a multivolume-tape data file:

- The sequence number of the new data file must be the sequence number of the last data file on the multivolume tape data file plus 1.

In the Figure about Data File Sequence the sequence number of FILEC must be 3, even though there are only two files on the last volume.

Always specify the location of a data file on tape with the SEQNBR parameter. The information specified on the LABEL parameter verifies that you found the correct data file. However, this verification only occurs after the system locates the file that is specified with the SEQNBR parameter. You cannot locate a data file on tape by label name. If you use the Check Tape (CHKTAP) command, the sequence number of the data file returns in the completion message.

You may use some special values in place of an actual sequence number:
- *NEXT: The system processes the next sequential data file on the tape. If you position the tape prior to the first data file, the system processes the first data file on the tape. *NEXT is useful for applications that need to read all data files on a tape. The system uses this value for tape device files that are used to read from tape. The system issues an error message when you use a tape device file to write to a tape and you specify *NEXT.
- *END: The system writes the data file to the end of the tape. The system uses this value in tape device files that are used to write to tape. The system issues an error message when you use a tape device file to read from a tape and you specify *END.

The SEQNBR parameter for an output file, for which EXTEND(*NO) is specified, must be one of the following:
- SEQNBR(1). This overwrites the first data file on the volume, regardless of the sequence number in the labels of the first data file already on the volume.

- A value of 1 greater than the value for a data file that already exists on the volume. This either overwrites an existing data file on the volume or adds a data file at the end of the volume.

    Note: If the tape device is a 1/4-inch or 8-mm cartridge device, the program will not overwrite the existing files.

- *END.

LABEL

The LABEL parameter specifies the data file label on the tape.

The information specified on the LABEL parameter is used for new labels created for an output file for which EXTEND(*NO) is specified. It is also used for an I/O file for which EXTEND(*YES) is specified to verify that the correct file is processed.

FILETYPE

The file type of the file to process. The value should be *DATA for a data physical file and *SRC for a source physical file. The program allows this parameter only on the CRTTAPF command.

RCDLEN

The parameter RCDLEN specifies the length of records that are used by a program using this device file. If *CALC is specified, the system attempts to calculate record length from the file header labels. The maximum record length is 32 767 bytes for fixed-length or undefined format records, and 32 759 for variable-length or spanned records. Fixed-length and undefined format output records cannot be less than 18 bytes in length.

BLKLEN

The BLKLEN parameter specifies the data block length that transfers on each I/O operation. If *CALC is specified, the system attempts to calculate block length from the file header labels. The block length must be between 18 and 524 288 bytes.

RCDBLKFMT

The RCDBLKFMT parameter specifies the format of the I/O records and blocks. Records can be:
- D-type ASCII, deblocks (*D)
- D-type ASCII, blocked (*DB)
- Fixed-length, deblocked (*F)
- Fixed-length, blocked (*FB)
- Variable-length, deblocked, unspanned (*V)
- Variable-length, blocked, unspanned (*VB)
- Variable-length, deblocked, spanned (*VS)
- Variable-length, blocked, spanned (*VBS)
- Undefined format (variable length) (*U)

The record length, block length, and record block format may not need to be specified for standard-labeled I/O tape data files specified as EXTEND(*YES). The system can take this information from the tape labels.If the program specifies a block length or record block format that does not match tape label specifications in the tape label:
- The system then assumes the tape label specification.

If the record length specified in the program does not match the length of the data:

- The system then truncates or pads the data to the length specified in the program.

EXTEND

New records may be added to the end of the data file on the tape by specifying the EXTEND parameter. The destruction of all remaining data files occurs if the data file is not the last data file on the tape. The destruction of all remaining data files also occurs when you over-write an existing data file. The extension uses the label specifications for record and block length that are specified in the label. EXTEND is valid only for 1/2-inch tape devices.

By specifying EXTEND(*YES *CHECK) the system checks the expiration date of the first data file following the data file being extended.

DENSITY

The system records, in the same density, all data files on a volume. You use the DENSITY parameter only to set the output volume density when you create the first data file on a volume that is not labeled. You use the volume label on a labeled tape to determine the density format. For valid values, see the *CL Reference (Abridged)*, SC41-5722-03 .

COMPACT

Allows the user to control device data compaction for output files. If you do not want to use data compaction to, specify *NO on the COMPACT parameter. If you specify *DEVD and the device does not support data compaction, the system ignores this parameter.

CODE The CODE parameter specifies the character code (EBCDIC or ASCII) for the data that is not labeled. For standard label tapes, the volume label determines the character code. The system writes ASCII Interchange code when the character code is ASCII. The data conforms to the "American National Standard" X3.27-1978, "Magnetic Tape, and File Structure for Information Interchange".

CRTDATE

The CRTDATE parameter specifies the creation date of an input data file on a labeled tape. The system sends a message to the system operator if the creation date on the tape does not match the date in the file description.

EXPDATE

The EXPDATE parameter specifies the expiration date of an output data file on a labeled tape. The program cannot write-over the data file until the date has expired. The program considers the data file protected.

The program may create an output data file instead of extending an existing data file. When this occurs, the system compares the expiration date of the new data file to the date of the file which precedes it on the volume. If the expiration date of the new data file is later than the file preceding it on the tape, the program sends an inquiry message (CPA4036). The system operator can choose one of the following operations:

- The creation of the data file.
- Load a new tape and try again
- Allow the program to end processing

Note: Creating the data file could produce a volume for which CHECK(*FIRST) on the INZTAP command is unreliable.

If you do not want the data file to be written over, specify *PERM on the
EXPDATE parameter.

ENDOPT

The ENDOPT parameter specifies where to position the magnetic tape
when the program closes the tape device file. The program:

- Rewinds the magnetic tape to the load point.
- Leaves the magnetic tape as it is.
- Unloads the magnetic tape.

When you use a multivolume-tape data file and specify ENDOPT(*LEAVE),
you must place the first volume on the first tape device specified in the DEV
parameter. (The exception to this is for a read backward, in which case you
must place the last volume on the first tape device specified). If a user
opens the data file again, with the same device list, and leaves the tape on
a different tape device:

- Place the tape volume on the first tape device that is specified in the
  DEV parameter before you open the next data file on that tape reel.

Note the following restriction when using *LEAVE processing with tape
media libraries. *LEAVE processing restricts the use of the resource that
has the current cartridge mounted to that same cartridge. A resource
allocation timeout will occur if both of the following conditions exist:

- The device is the only resource available to the media library.
- The program issues a request to use a different cartridge.

The resource will remain unavailable to the program until:

- The program issues a command to rewind or unload the cartridge.
- The job that left the cartridge in *LEAVE processing ends.

USRLBLPGM

The command supports user header and trailer labels through the use of
the USRLBLPGM parameter. USRLBLPGM identifies the user program that
is used during open and close processing. See "Processing User Labels" on
page 28 for more information.

BUFOFSET

The buffer offset length for an ASCII file is specified using the BUFOFSET
parameter. You can specify a buffer offset length for any ASCII input data
file. You can specify a buffer offset value of

- *BLKDSC for an input or output ASCII formats *D file
- *BLKDSC for an input or output ASCII formats *DB file

to process a block with 4-digit block descriptors.

For additional tape information, and information about using tape for save and
restore operations, see the *Backup and Recovery*, SC41-5304-03 book. Table 1 on
page 15 lists parameters that apply to magnetic tape and where to specify the
parameters. The *CL Reference (Abridged)*, SC41-5722-03 manual contains detailed
information about how to specify these parameters on the CRTTAPF, CHGTAPF,
and OVRTAPF commands.

Table 1. Tape Device File Parameters

CL Parameter	Description	Specified on CRTTAPF Command	Specified on OVRTAPF Command	Specified in HLL Programs
FILE	File name	Qualified file name	File name	ILE RPG, COBOL, BASIC, PL/I, or ILE C programming languages
DEV	Device name	*NONE or list of device names	List of device names	
VOL	Volume	*NONE or list of volume identifiers	*NONE or list of volume identifiers	
REELS	Volume label type	*SL, *NL, *NS, *BLP, or *LTM	*SL, *NL, *NS, *BLP, or *LTM	
REELS	Number of labeled tapes	Number of reels	Number of reels	
SEQNBR	Sequence number	*NEXT, *END, or sequence number of file	*NEXT, *END, or sequence number of file	
LABEL	Label	*NONE or file label	*NONE or file label	BASIC
FILETYPE	File type	*DATA or *SRC		
RCDLEN	Record length	*CALC or record length	*CALC or record length	ILE RPG, COBOL, BASIC, PL/I, or ILE C programming languages
BLKLEN	Block length	*CALC or block length	*CALC or block length	COBOL programming language
BUFOFSET	Buffer offset	*BLKDSC or buffer offset	*BLKDSC or buffer offset	
RCDBLKFMT	Record block format	*F, *FB, *V, *VB, *D, *DB, *VS, *VBS, or *U	*F, *FB, *V, *VB, *D, *DB, *VS, *VBS, or *U	COBOL, ILE C programming languages
EXTEND	Extend	*NO, *YES *CHECK, or *YES *NOCHECK	*NO, *YES *CHECK, or *YES *NOCHECK	COBOL, ILE C programming languages
DENSITY	Density	See the *CL Reference (Abridged)*.	See the *CL Reference (Abridged)*.	
COMPACT	Data compaction	*DEVD or *NO	*DEVD or *NO	
CODE	Character code	*EBCDIC or *ASCII	*EBCDIC or *ASCII	COBOL programming language
CRTDATE	Creation date	*NONE or date	*NONE or date	
EXPDATE	Expiration date	*NONE, date, or *PERM	*NONE, date, or *PERM	
ENDOPT	End option	*REWIND, *LEAVE or *UNLOAD	*REWIND, *LEAVE or *UNLOAD	COBOL programming language
USRLBLPGM	User label program	*NONE or qualified program name	*NONE or qualified program name	
IGCDTA	Double-byte data	N/A	*NO or *YES	
WAITFILE	File wait time	*IMMED, *CLS, or number of seconds	*IMMED, *CLS, or number of seconds	
SHARE	Shared file	*NO or *YES	*NO or *YES	
AUT	Authority	*LIBCRTAUT, *CHANGE, *ALL, *USE, *EXCLUDE, or authorization list name	N/A	
REPLACE	Replace existing file	*YES or *NO	N/A	
TEXT	Text	*BLANK or text	N/A	

Using Tape Device Files in High-Level Language Programs

A program-described device file can access a magnetic tape device. To use a tape device file with a program, you must either specify the tape file name in the program or use an Override with Tape File (OVRTAPF) command. The high-level language determines what tape parameters to use in the program.

Open Processing for Tape Device Files

The following considerations apply to opening tape device files:

- When the program opens a tape device file, the system merges any parameters that are specified in the file with the parameters that are specified in the program. The system then merges these parameters with the parameters that are specified on an OVRTAPF command.

- Specify the device names when you open the tape device file. If you specify DEV(*NONE) in the tape file, you must specify one or more device names on an OVRTAPF command. You can specify as many as four device names for a single tape device file (depending on how many magnetic tape devices you have).

  The record length, block length, record-block-format, and buffer offset (for an ASCII file) always return to the program in the data management open feedback area. They return in the format in which they are written in the HDR2 file header label. This information is available regardless of the type of label processing that is used for the file.

- The following data files support the read-backward operation:
  - fixed (*F)
  - fixed block (*FB)
  - and undefined format (*U)

  for both single volume tape data files and multivolume-tape data files. You request a read backward operation through a high-level language when you open the file. An escape message is signaled if a read-backward operation is attempted for variable-length (spanned or not spanned) or source records.

  Note: The following tape devices do not have read-backward capabilities: Use the Retrieve Device Capabilities (QTARDCAP) to determine the device capabilities of your device. See the *System API Programming*, SC41-5800-00 for information about QTARDCAP.

  - 9348 tape unit
  - 8-mm cartridge device
  - Some 1/4-inch cartridge devices

  Reading a data file backwards, where you specify the device and volume list, you must position the volumes on the device in reverse order. For example, a device file with DEV(QTAPE1 QTAPE2) VOL(VOL01 VOL02 VOL03) expects VOL03 on QTAPE1, VOL02 on QTAPE2, and then VOL01 on QTAPE1.

  For a read-backward operation the end-of-file condition occurs if the system recognizes the first volume of the data file from the header labels for:
  - Standard label processing (*SL)
  - Bypass label processing (*BLP)

  If the system does not recognize the header labels for the first volume of the data file or if this is a *BLP file:

- The system signals the end-of-file condition when:
  - The system processes the specified number of reels.
  - The system processes the number of identifiers on the VOL parameter.
- Some high-level languages allow you to specify where to position the tape when the program opens an input tape device file. This indicates whether you process the tape in a forward or the backward direction. The rules that determine the first volume of a data file from the labels are:
  - HDR1 labels multivolume sequence field = 1 (ASCII or EBCDIC with no HDR2 label)

    or
  - HDR2 label volume switch indicator field = 0 (EBCDIC)

  The program specifies the record length according to the information that is shown in Table 2 on page 18.
- For source files, the record length used to determine the block length is the actual data length, not the data length plus 12 bytes (for sequence number and date).
- You must supply either a RCDLEN or BLKLEN parameter value for unspanned, deblock records (*F, *V, *D, *U).
- You must supply both RCDLEN and BLKLEN parameter values for spanned or blocked records (*FB, *VB, *DB, *VS, *VBS).
- When the file type specified in the tape device file is a source file:
  - The system appends a date and sequence number to each record on input operations. The date field is always 0.
  - The system removes the date and sequence number from each record on output operations

  The program can check (if the high-level language you are using allows it) to determine if the input or output file is a source file. The record length must include 12 bytes for the date and sequence number. The block length and record length ratio remains the same for a source block and data record minus the 12 bytes allocated to source files. For example, if the actual data record length is 80 the record length becomes 92 for a source file. The block length remains unchanged.
- To process input files using standard labels the system will always use the block length in the file label. The device file block length is ignored.
- Variable-length (spanned or unspanned) records and undefined format records can be used for output files. If your high-level language does not support variable-length records, then all records for an output tape device file that uses variable format are maximum length.
- Specify the sequence number so you can find the data file. You cannot locate tape data files by label name.
- When you specify both the VOL and REELS parameters the REELS parameter is ignored. If you want to use the REELS parameter (number of reels) to limit the number of input volumes that is processed, specify *NONE for the VOL parameter.

Table 2. Specifying Record Lengths by Record and Format Type

Record and Format Type	Minimum Record Length for *DATA	Minimum Record Length for *SRC	Maximum Record Length for *DATA	Maximum Record Length for *SRC	Block Length
Fixed blocked, *F, *FB, *U	18	30	32 767	32 767	Multiple of *DATA record length
Variable unblocked, *V	1	13	32 759 (See Note)	32 767	Equal to maximum *DATA record length plus 8
D-type ASCII unblocked, *D	1	13	9 995 (See Note)	10 007 (See Note)	Equal to maximum *DATA record length plus 4, plus buffer offset
Variable blocked, *VB	1	13	32 759	32 767	At least maximum *DATA record length plus 8
D-type ASCII blocked, *DB	1	13	9 995 (See Note)	10 007 (See Note)	At least maximum *DATA record length plus 4, plus buffer offset
*VS, *VBS	1	13	32 759	32 759	

: This is the maximum record length for a record being written to a tape. Input records can be padded to 32 767.

Input/Output Processing for Tape

The following considerations apply to I/O operations that are performed on data files:

Read and Write Considerations

- Specify record lengths in your program when writing variable-length records (*D, *DB, *V, *VB, *VS, *VBS, or *U specified on the RCDBLKFMT parameter on the CRTTAPF command).

  If the maximum record length on tape is shorter than the output record length (because of overrides or existing file labels):

  – The system truncates the records to the maximum length that is allowed.
  – The system sends a diagnostic message when you open a device file to indicate that output records may be truncated.

- If the program specifies a record length that differs from the actual length of the data, the system pads or truncates the data to match the program specification.

Read Considerations

- If the system does not find an end-of-file label, processing will continue until the system uses all the specified volume identifiers. If the program specifies VOL(*NONE , then the system processes the tapes until reaching the specified number of reels (REELS parameter). The system sends a message CPA5230 to the system operator message queue when all the identifiers in the VOL list are processed. When you receive this message, you can do the following:

  – Cancel processing of the device file immediately.The system will close the device file.
  – Continue: To process other volumes.

- When the system reads a tape block that is not a valid length, the system sends a CPF5036 notify message. If your high-level language reports this condition to

your program, it can continue to process by reading another record. When you continue this way, the system skips the block that is not valid so your program does not receive any records from this block.

Force-End-of-Data Considerations

The force-end-of-data function is valid for both input and output. The force-end-of-data function for an output file forces the system to write all buffered records to tape. When this occurs, the system may write a short block on the volume. The force-end-of-data function for an input file positions the tape at the last volume for the file and signals end-of-file to the using program.

Force-End-of-Volume Considerations

The force-end-of-volume function is valid for both input and output files. It causes volumes to switch immediately, or signals end-of-file if there are no more continuation volumes for an input file.

Close Processing for Tape

When you close a tape device file, the system performs one of several functions according to what you specify in the tape device file. The system uses the ENDOPT parameter on a CRTTAPF, CHGTAPF, or an OVRTAPF command. The system can perform the functions that are listed below:

- The system rewinds the tape.
- The system leaves the tape as it is.
- The system rewinds and unloads the tape to remove it from the magnetic tape device.

If the tape operation ends abnormally:

- The position of the tape when you close the device file may be unchanged.
- The system may rewind the tape regardless of specified program instructions.

Understanding Records, Blocks, and Formats

This topic provides information to help you understand records, blocks, and record formats of the tapes.

- Records

  Records are logical mappings of the data on a tape. It usually maps directly to the records in a database file.
- Blocks

  Blocks are physical units of data on a tape. Blocks can include a record, part of a record, or multiple records.
- Record block format

  The record block format allows the system and user to interpret the data on a tape.

To understand records, blocks, and record block formats you must know a few key terms:

- Fixed length

  The blocks on a tape have an exact (or fixed) length.
- Variable length.

The blocks on a tape have a variable length. The block contains a header with the length of the block. Each block in a file may or may not have the same length.

- Undefined length

  The blocks on a tape have no defined length, each block can be different, and the program application interprets each block correctly.

- Blocked records

  Blocking is the process of grouping records into blocks before the system writes records on a volume. A block consists of one or more logical records. Blocking conserves storage space on a volume by reducing the number of interblock gaps in the data set. This increases processing efficiency by reducing the number of I/O operations that are required to process the data set.

- Deblocked records

  One record exists per block.

- Spanned records

  The system splits a single record (spans) into two different data blocks and writes them on the tape.

- Unspanned records

  Each record is contained within one data block.

- Interblock gap

  The physical gap on tape between two data blocks.

- Block Descriptor Word (BDW)

  One or more logical records or record segments follow a block descriptor word (BDW) in a variable length block.

- Record Descriptor Word (RDW)

  Data follows a record descriptor word (RDW) in a variable length logical record. The RDW describes the record.

- Record Segment

  Since spanned records occupy more than one block, each part of the record is a record segment.
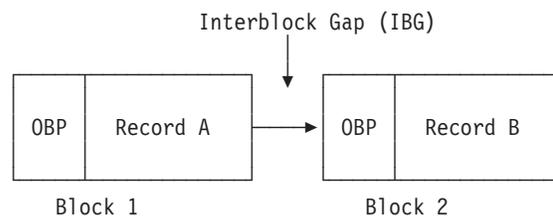
- Segment Descriptor Word (SDW)

  Data follows a segment descriptor word (SDW) in each record segment. The SDW, similar to the RDW, describes the record segment.

In sorting out these terms, the program supports and translates certain combinations into the following record block formats.
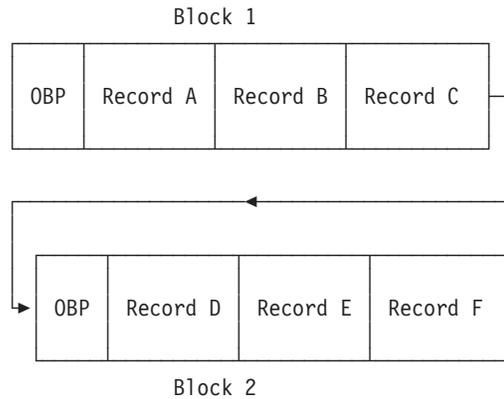
Example - Record format *D

D-type Code ASCII, deblocks (*D).

```
                        Interblock Gap (IBG)
                                 |
                                 v
  +------+-----------+      +------+-----------+
  | OBP  | Record A  | ---> | OBP  | Record B  |
  +------+-----------+      +------+-----------+

       Block 1                   Block 2
```

OBP is an optional block prefix. The buffer offset (BUFOFSET) parameter, for the tape file, specifies the optional block prefix that can vary in length from 00 to 99. The OBP length must be constant for all blocks in the file. Each record can also contain an optional control character.
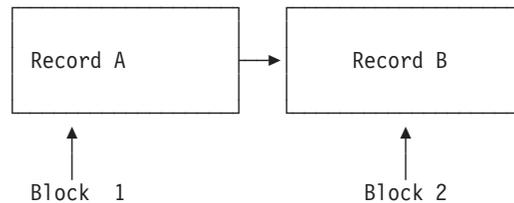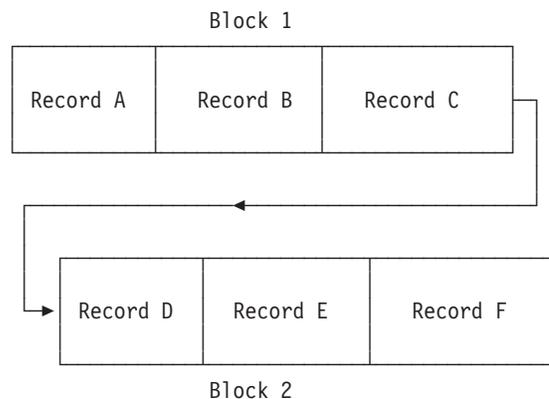
Example - Record format *DB

D-type code ASCII, blocked (*DB).

```
                      Block 1

  +-------+----------+----------+----------+
  | OBP   | Record A | Record B | Record C |
  +-------+----------+----------+----------+


  +-------+----------+----------+----------+
  | OBP   | Record D | Record E | Record F |
  +-------+----------+----------+----------+

                      Block 2
```

OBP is an optional block prefix. The buffer offset (BUFOFSET) parameter, for the tape file, specifies the optional block prefix that can vary in length from 00 to 99. The OBP length must be constant for all blocks in the file. Each record can also contain an optional control character.

Example - Record format *F
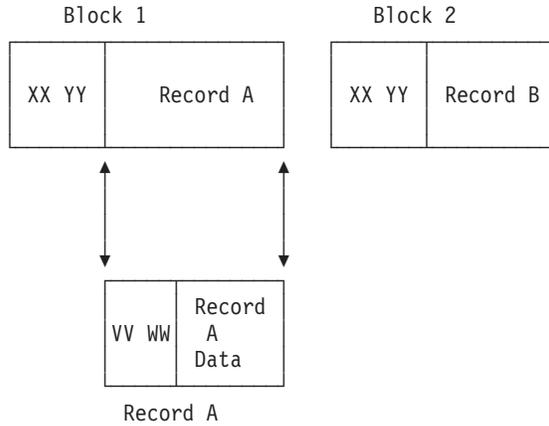
Fixed-length, deblocked (*F).

```
  +-------------+        +-------------+
  |             |        |             |
  |  Record A   | -----> |  Record B   |
  |             |        |             |
  +-------------+        +-------------+

     Block  1               Block 2
```

Example - Record format *FB

Fixed-length, blocked (*FB).

```
                      Block 1

  +----------+----------+----------+
  | Record A | Record B | Record C |
  +----------+----------+----------+


  +----------+----------+----------+
  | Record D | Record E | Record F |
  +----------+----------+----------+

                      Block 2
```

Example - Record format *V

Variable-length, deblocked, unspanned (*V).

```
        Block 1                    Block 2
 +--------+----------+       +--------+----------+
 |        |          |       |        |          |
 | XX YY  | Record A |       | XX YY  | Record B |
 |        |          |       |        |          |
 +--------+----------+       +--------+----------+
           ^        ^
           |        |
           |        |
           v        v
       +-------+--------+
       |       | Record |
       | VV WW |   A    |
       |       |  Data  |
       +-------+--------+
            Record A
```

XX and YY make up the BDW. XX is the length of the record plus the length of the BDW (4 bytes). YY are currently reserved fields and must be 00. XX should be the actual length of the data block that is written.
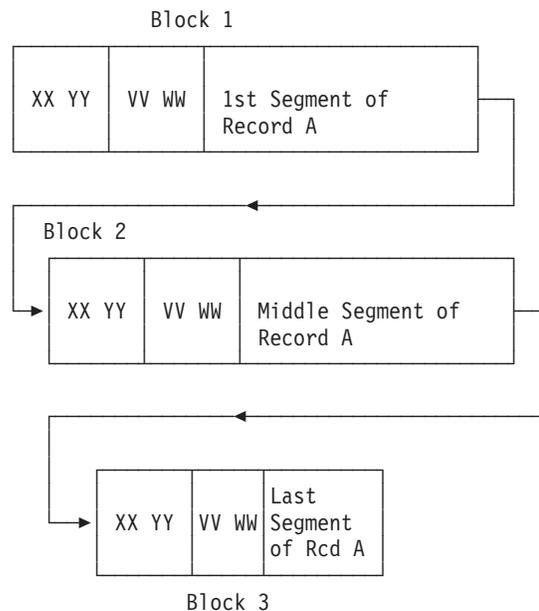
Record A has its own mapping by including the RDW. VV and WW make up the RDW. VV is the length of the record plus the length of the RDW (4 bytes). WW are currently reserved fields and must be 00. VV should be the actual block length (the value in XX) minus 4 bytes (size of BDW). OS/400 will pad blocks to make the 18 byte block limit. This occurs if the record data is less than 10 bytes (10 bytes plus 8 bytes of header information is the 18 byte block limit).

Example - Record format *VB

Variable-length, blocked, unspanned (*VB).

```
            +-------+--------+
            |       | Record |
            | VV WW |   A    |
            |       |  Data  |
            +-------+--------+
                |        |
                |        |
                |        |
 +-------+----------+----------+
 |       |          |          |
 | XX YY | Record A | Record B |
 |       |          |          |
 +-------+----------+----------+
    ^           |        |
    |           |        |
    |           |        |
 Block 1    +-------+--------+
            |       | Record |
            | VV WW |   B    |
            |       |  Data  |
            +-------+--------+
```
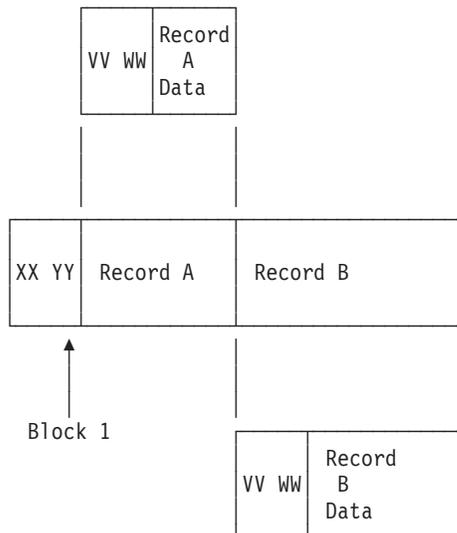
XX and YY make up the BDW. XX is the length of all records plus the length of the BDW (4 bytes). YY are currently reserved fields and must be 00. XX should be the actual length of the data block that is written.

```
                      Record A
                         |
                         v
        +--------+------------------+
        |        |                  |
        | VV  WW |  Record A data    |
        |        |                  |
        +--------+------------------+
```
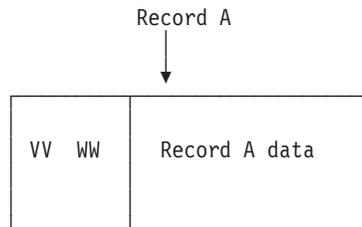
Record A has its own mapping by including the RDW. VV and WW make up the RDW. VV is the length of the record plus the length of the RDW (4 bytes). WW are currently reserved fields and must be 00. VV should be the actual length of record data A plus the length of the RDW (4 bytes). The actual length of the data block equals the sum of:

- The VV value for ALL records.
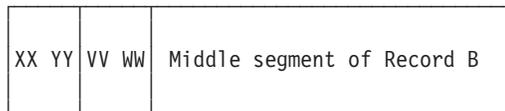- The length of the BDW (4 bytes).

Example - Record format *VS

Variable-length, deblocked, spanned (*VS).

```
              Block 1
  +-------+-------+----------------+
  |       |       |                |
  | XX YY | VV WW | 1st Segment of  |
  |       |       | Record A        |
  +-------+-------+----------------+

    Block 2
    +-------+-------+--------------------+
    |       |       |                    |
    | XX YY | VV WW | Middle Segment of   |
    |       |       | Record A            |
    +-------+-------+--------------------+

      +-------+-------+----------+
      |       |       | Last      |
      | XX YY | VV WW | Segment   |
      |       |       | of Rcd A  |
      +-------+-------+----------+
              Block 3
```
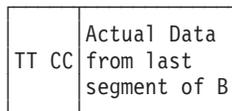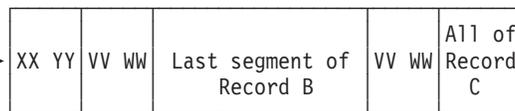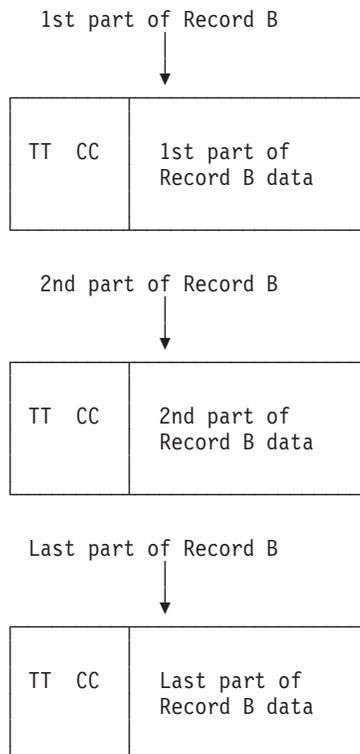
XX and YY make up the BDW for each data block. XX is the length of all records in each block plus the length of the BDW (4 bytes). YY are currently reserved fields and must be 00. XX should be the actual length of the data block that is written.
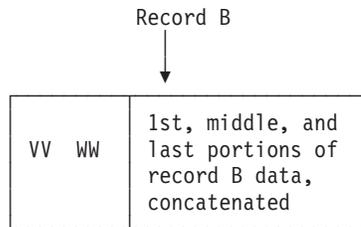
Note that logical record A spans over three actual data blocks on the tape.

By breaking down each piece of record A, you can see that each segment has its own characteristics.

```
+-------+-----------------+
|       | 1st part of     |
| TT CC | Record A        |
|       | Data            |
+-------+-----------------+
```

```
+-------+-----------------+
|       | Middle part of  |
| TT CC | Record data A   |
|       |                 |
+-------+-----------------+
```

```
+-------+-----------------+
|       | Last part       |
| TT CC | of              |
|       | Record A        |
+-------+-----------------+
```

Note that the actual mapping of the entire block 3 is the following:

```
+-------+-------+-------+-----------------+
|       |       |       | Last part of    |
| XX YY | VV WW | TT CC | Record A        |
|       |       |       |                 |
+-------+-------+-------+-----------------+
```

Each segment of record A has its own by mapping by including the SDW. TT and CC make up the SDW. TT is the length of the record plus the length of the SDW (4 bytes). CC is the segment control character. The first byte of the CC defines which part of the record the segment is. The values of the control character can be:

- 00 binary - Complete logical record
- 01 binary - First segment of a multi-segment record
- 10 binary - Last segment of a multi-segment record
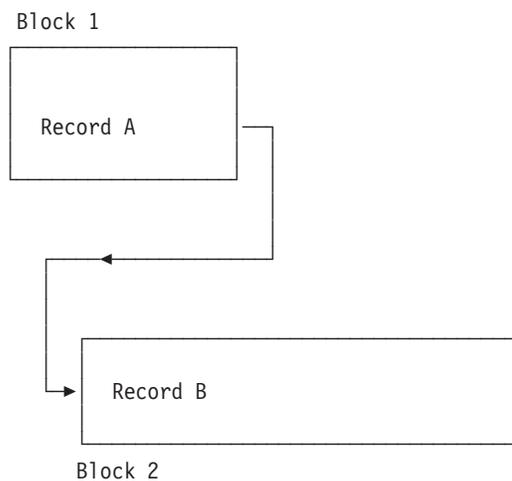- 11 binary - Middle segment of a multi-segment record

The second byte of the control character is reserved, and should be 0. TT should be the actual length of record data segment plus the length of the SDW (4 bytes).

From a user's point of view, a logical view of record A is the same as other records defined above.

```
                Record A
                   |
                   v
        +-------+-----------------+
        |       | 1st, middle, and|
        | VV WW | last portions of|
        |       | record A data,  |
        |       | concatenated    |
        +-------+-----------------+
```
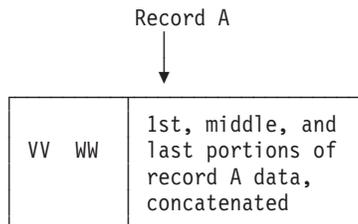
Example - Record format *VBS

Variable-length, blocked, spanned (*VBS).

Block 1
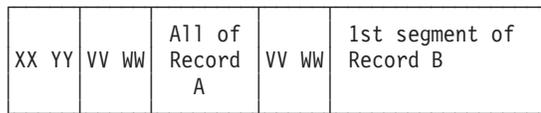
XX YY	VV WW	All of Record A	VV WW	1st segment of Record B

Block 2

XX YY	VV WW	Middle segment of Record B

Block 3

XX YY	VV WW	Last segment of Record B	VV WW	All of Record C

TT CC	Actual Data from last segment of B

```
     1st part of Record B

               │
               ▼
 ┌──────────┬─────────────────┐
 │          │                 │
 │ TT  CC   │ 1st part of     │
 │          │ Record B data   │
 │          │                 │
 └──────────┴─────────────────┘


     2nd part of Record B

               │
               ▼
 ┌──────────┬─────────────────┐
 │          │                 │
 │ TT  CC   │ 2nd part of     │
 │          │ Record B data   │
 │          │                 │
 └──────────┴─────────────────┘


     Last part of Record B

               │
               ▼
 ┌──────────┬─────────────────┐
 │          │                 │
 │ TT  CC   │ Last part of    │
 │          │ Record B data   │
 │          │                 │
 └──────────┴─────────────────┘
```

XX and YY make up the BDW for each data block. XX is the length of all records in each block plus the length of the BDW (4 bytes). YY are currently reserved fields and must be 00. XX should be the actual length of the data block that is written.

Note that logical record B spans across three actual data blocks on the tape.

Each segment of record B has its own by mapping by including the SDW. TT and CC make up the SDW. TT is the length of the record plus the length of the SDW (4 bytes). CC is the segment control character. The first byte of the CC defines which part of the record the segment is. The values of the control character can be:

- 00 binary - Complete logical record
- 01 binary - First segment of a multi-segment record
- 10 binary - Last segment of a multi-segment record
- 11 binary - Middle segment of a multi-segment record

The second byte of the control character is reserved, and should be 0. TT should be the actual length of record data segment plus the length of the SDW (4 bytes).

From a user's point of view, a logical view of record B is the the same as other records defined above.

```
                       Record B
                          │
                          ▼
        ┌───────┬─────────────────────┐
        │       │ 1st, middle, and    │
        │ VV  WW│ last portions of    │
        │       │ record B data,      │
        │       │ concatenated        │
        └───────┴─────────────────────┘
```

Example - Record format *U

Undefined format (variable length) (*U.)

```
   Block 1
 ┌─────────────────────┐
 │                     │
 │                     │
 │ Record A        ────┼──┐
 │                     │  │
 │                     │  │
 └─────────────────────┘  │
         ┌────────────────┘
         │
         ▼
    ┌──────────────────────────────────┐
    │                                  │
 ┌─▶│ Record B                         │
 │  │                                  │
 │  └──────────────────────────────────┘
 │     Block 2
```

Undefined records do not have a standard definition, and the contents of Records A
and B are defined by the application. The application may define its own descriptors
with lengths similar to BDWs.

Handling Tape Processing Errors and Damaged Tapes

When the system does not write tape marks or labels at the end of the data file,
damage to the data files can occur. This can happen due to an error condition or a
system failure. If this happens when writing a file to a 1/2-inch tape device, the
following occurs when you try to read the data file:

- The new data and existing data may appear concatenated when processed for
  input. When the tape has labels the system sends an error message to the
  system operator when the program reads the trailer labels. The system does not
  detect an error for unlabeled tape.
- If the new data and the existing data do not appear concatenated, then the
  system sends an error message to the system operator.
- If the tape you use as input:
  - Is a tape that does not contain existing data or tape marks beyond the
    location of the last data block.
  - Is a new tape.
  - Is a completely erased tape.

Then the tape advances off the end of the reel.

Note: Whenever you close an output file:

- The system attempts to write an end-of-tape marker and label at the end of the file.
- If the system cannot write closing tape marks and labels, the system sends a message to your job log.

When an error occurs when you save data by using a SAVxxx command, the system prompts the operator to load another tape or cancel the job. If the operator loads another tape volume:

- The system repositions the tape a number of blocks before the error occurred.
- The system writes end-of-volume labels.

The job then continues to write data starting with the first block of data that was overwritten with end-of-volume labels.

If damage occurs while using a cartridge tape device, and you encounter a blank tape, the program sends an error message to the system operator.

Processing User Labels

The system uses the USRLBLPGM parameter to specify the name of a program used to process user header and trailer labels. This parameter is not valid for save and restore functions.

The system calls the program, which the USRLBLPGM parameter specifies, for open and close processing, to process every label. The system calls the program one additional time to tell the program that there are no more labels.

Figure 5 on page 29 shows the format of a tape with user labels. At open, the system will call the user-label program three times for processing the labels in the diagram. Those calls are for UHL1, UHL2, and a final time to indicate completion. At close, the system will call the user-label program three more times.

The system passes three variables to the program. The program variables have the following lengths:

- Parameter 1: 80 characters
- Parameter 2: 1 character
- Parameter 3: 244 characters

The following list shows the content the variables pass to the user label program:

Parameter 1

Position 1 - 80

User header or trailer label

- For output files the program sets this variable to the next user label to write to tape.
- For input files written to the system it sets this variable to the user label that is most recently read from the tape.

Parameter 2

Position 1

End-of-labels indication

Parameter 2 contains a character 0 or 1, that indicates whether or not the label read is the last label read. For output files, the user label program sets the value. For input files, the system sets the value.

- 0 indicates that Parameter 1 contains a label.
- 1 indicates that Parameter 1 does not contain a label. There are no labels to process.

VOL	HDR1	HDR2	UHL1	UHL2	TM	Data

	TM	TRL1	TRL2	UTL1	UTL2	TM	TM

RV2W703-0
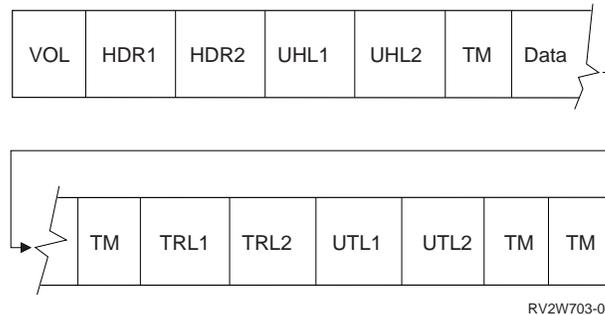
Figure 5. Tape with User Labels

Parameter 3

Position 1- 80
> Current volume label

Position 81-160
> The last processed HDR1 or TRL1 label.

Position 161-240
> The last processed HDR2 or TRL2 label.

Position 241-242
> User label number
> - Output files: The number of the next user label to be written in the current header or trailer group.
> - Input files: The total count of user labels read in the current header or trailer group.

Position 243
> Open file option
>
> The open file option field contains a character that indicates whether the file is open for input or for output.
> - I Indicates that the file is an input file.
> - O Indicates that the file is an output file.

Position 244
> Expect labels
>
> The expect labels field contains an integer indicating whether the call is returning or expecting labels from the user program.
> - 0 indicates that the user program is returning labels.
> - 1 indicates that the user program is expecting labels.

Other Tape Support Commands

The following CL commands provide tape support functions.

CHKTAP

Check Tape: The check tape command searches a tape volume for a specific volume identifier or field label.

The following command checks a tape volume, that is loaded on device TAP01, for the location of a file named TEST at sequence number 3.

```
CHKTAP  DEV(TAP01) SEQNBR(3) LABEL(TEST)
```

DMPTAP

Dump Tape: This command dumps label information, data blocks, or both, from a tape with or without a label.

The following command dumps the second file, including all label information, from a tape volume that is loaded on device TAP02.

```
DMPTAP  DEV(TAP02) SEQNBR(2) TYPE(*ALL)
```

DSPTAP

Display Tape: The command displays:
- Volume and file information for labeled volumes.
- Saved object information.
- Volume type (*NL or *LTM) and density for volumes without labels.

The following command displays the volume and file label information of a tape volume loaded on device TAP01.

```
DSPTAP  DEV(TAP01)
```

ADDTAPCTG

Add Tape Cartridge: The command adds the specified cartridge to a valid category for media library automation.

The following command adds the cartridge PSH123 to the *SHARE400 category on media library device TAPMLB01.

```
ADDTAPCTG DEV(TAPMLB01) CTG(PSH123)
CGY(*SHARE400)
```

RMVTAPCTG

Remove Tape Cartridge: The command removes the specified cartridge to the convenience station or high capacity station of the media library device.

The following command removes the cartridge PSH123 to the convenience station on the media library device TAPMLB01.

```
RMVTAPCTG DEV(TAPMLB01) CTG(PSH123)
```

CHGTAPCTG

Change Tape Cartridge: The command changes the category attribute of the specified cartridge.

The following command changes the category of cartridge PSH123 to the *NOSHARE category on media library device TAPMLB01.

```
CHGTAPCTG DEV(TAPMLB01)
CTG(PSH123) CGY(*NOSHARE)
```

DSPTAPCTG

Display Tape Cartridge: The command displays the attributes and information about the specified cartridge.

The following command displays the category and other information about cartridge identifier PSH123 on media library device TAPMLB01.

```
DSPTAPCTG DEV(TAPMLB01) CTG(PSH123)
```

WRKTAPCTG

Work with Tape Cartridge: The command works with the specified cartridges in the media library device.

The following command line will display a work panel for all tape cartridges in media library device TAPMLB01.

```
WRKTAPCTG DEV(TAPMLB01)
```

SETTAPCGY

Set Tape Category: The command mounts a category on the specified media library device. Mounting a category refers to creating a sequential loader (loads tape cartridges in order) in the order they were placed into the specified category. The mounted resource is used by specifying VOL(*MOUNTED) on any tape related command: CPYTOTAP, SAVLIB, DSPTAP, INZTAP, and so on.

The following command mounts the category FRIDAY to the media library device TAPMLB01.

```
SETTAPCGY DEV(TAPMLB01)
OPTION(*MOUNTED) CGY(FRIDAY)
```

CRTTAPCGY

Create Tape Category: The command creates a user defined category to assign tape cartridges.

The following command creates a user defined category called SCRATCH.

```
CRTTAPCGY CGY(SCRATCH)
```

DLTTAPCGY

Delete Tape Category: The command deletes a user defined category that was previously created. The category cannot contain cartridges.

The following command deletes a user defined category called JOE.

```
DLTTAPCGY CGY(JOE)
```

DSPTAPCGY

Display Tape Category: The command displays the current list of valid tape categories.

The following command displays the existing categories.

```
DSPTAPCGY OUTPUT(*)
```

CFGDEVMLB

Configure Device Media Library: The command configures the connection between an existing media library device description and the communications resource that drives the robotics of the device.

The following command configures (connects) the media library device TAPMLB01 with the communication resource CMN01.

```
CFGDEVMLB DEV(TAPMLB01) ADPTTYPE(*RS232)
RSRCNAME(CMN01)
```

DSPLANMLB

Display LAN Media Library: The command displays the LAN information necessary to configure a library manager console.

The following command displays the LAN information about TRNLINE line description so configuration on the library manager console can be done for a *LAN media library device.

```
DSPLANMLB LIND(TRNLINE)
```

WRKMLBSTS

Work with Media Library Status: The command works with the status of the specified media library and its corresponding resources.

The following command works with all media libraries on the system and their corresponding tape resources.

```
WRKMLBSTS MLB(*ALL)
```

DUPTAP

Duplicate Tape: The command copies the contents from one tape volume to another.

The following command duplicates the tape volume mounted on device TAPE01 onto the tape volume mounted on device TAPE02.

```
DUPTAP FROMDEV(TAPE01) TODEV(TAPE02)
```

Record Formats

The header 2 label definition provides three kinds of record formats: fixed, variable, and undefined.

- F - Fixed length records
- V - Variable length records
- U - Undefined length records

Tape Management Systems (TMS)

Media and Storage Extensions (MSE) optional feature of OS/400 allows a user to write a tape management system, or purchase a tape management application from a provider. OS/400 recommends a product called Backup Recovery Media Services (BRMS/400). A tape management system is provided with exits through normal OS/400 tape processing. These exits allow a tape management system the ability to track tape volumes, protect tape volumes, provide scratch media, and an abundance of other features. Applications such as BRMS/400 can be purchased. The users can also use simple tape management systems to take advantage of one or many features provided through the MSE option. See the Tape Management Exit Program in the *System API Programming*, SC41-5800-00 for more information.

Understanding Device Capabilities

Retrieve Device Capabilities (QTARDCAP) is an API that can do the following:

- Return all device capabilities that are known by the system to the caller of the API.
- Provide information on whether the device supports hardware data compression (HDC) and so on.

For more information on QTARDCAP, see the *System API Programming*, SC41-5800-00.

Automation of Tape

Automated tape libraries (ATLs) or media library devices (MLBs) are the newest technology that enables customers to strive to full automation of backup and recovery. Full support of ATLs are in OS/400. Simply specify the device name and the cartridge identifier of the volume to be mounted. OS/400 will mount and complete the command issued. For more information about ATLs, see the *Automated Tape Library Planning and Management*, SC41-5309-02 book.

Archive and Recall in OS/400

Archive and Recall function is part of the Media and Storage Extensions (MSE) optional feature of OS/400. The Archive and Recall function can do the following:

- Enable archive and recall of the suspended objects in OS/400.
- Enable an exit program to restore a suspended object that is referred by a user without the user realizing that the object was not on the system (other than a slower execution).
- Be used in connection with automated tape libraries to fully automate the process.

For more information, see the Storage Extension Exit Program information in the *System API Programming*, SC41-5800-00.

Fast Access on Tape

Fast Access function is part of the Media and Storage Extensions (MSE) optional feature of OS/400. This function can Fast access a tape file. When opening a tape file for input, the Fast Access function can locate the tape file faster than the traditional sequential file search. When the tape file is created, a logical block identifier of the tape file is stored by a tape management system. The logical block is used when the tape file is accessed by the tape management system that supplies the logical block identifier of the tape file. For more information, see the Tape Management Exit Program information in the *System API Programming*, SC41-5800-00.

Considerations for Using the DUPTAP Command

To use the DUPTAP command your system must have two tape drives. In addition, these considerations apply:

- The density field in the file labels is updated to reflect the true density.
- Byte 80 on the volume label on a 6157 tape is reset from a Q to a blank.
- Tape duplication is not necessarily a volume to volume relation because of the different tape capacities.

  When duplicating a tape, the tape being copied to should have a larger capacity than the original tape (the one being copied from).

- Block length/size restrictions.

  If a tape is created by a save operation with the use optimum block (USEOPTBLK) parameter set to *YES, it may not be possible to duplicate it to another tape. For example, if USEOPTBLK(*YES) is specified for a save to a 3590 tape device, the resulting tape can only be duplicated to another device that can also support the larger block size.

- Concatenation has been a recently added function to the existing function of DUPTAP.

  Any number of standard label tapes can be duplicated to a single volume by specifying all volumes to be duplicated in the FROMVOL parameter. This feature has been added to allow duplicating to high capacity cartridges such as the 3590. For example, with this function, 40 3490 tape volumes can be duplicated to 1 3590 cartridge. The restriction that the 40 volumes be a multi-volume set is no longer required.

Performance Considerations for Tape

This section addresses performance concerns of the 3590 device. But it is not limited in scope to just the 3590. In fact, all of the high performance tape devices on the AS/400 can benefit from the following suggestions. Devices with high performance can reach data rate of 3 MB/s or more. The following describes the problems that may occur during the degradation of high performance tape devices. These problems may also occur in lower performing tape devices.

- Hardware Data Compression vs Software Data Compression

  Some of the older AS/400 attachments (Input/Output Processors) provide data compression in the data path hardware. This is referred to as Hardware Data Compression (HDC). Its purpose is to increase the data rate and capacity of the attached tape device. For compatibility with Input/Output Processors that do not provide HDC, the HDC algorithm is implemented in the system through Software Data Compression (SDC). SDC provides a performance increase for the entry level tape devices. For the high-end devices, SDC is a limiting factor to high performance. HDC and SDC are controlled by the data compression (DTACPR) parameter of the SAVE function commands in OS/400. If DTACPR(*DEV) is specified, then the outcome is also dependent on whether Improved Data Recording Capabilities (IDRC) is supported on the device.

  When DTACPR(*DEV)is specified, HDC is used if supported (only if IDRC is not supported). No SDC will occur. When DTACPR(*YES) is specified, HDC is used if supported. Otherwise SDC is used.

  DTACPR(*YES) should only be specified for a device attached to a type 2621 or 2644 Input/Output Processor. Specifying DTACPR(*YES) for devices attached to other types of Input/Output Processors will cause a major degradation in performance.

- Load balancing across the system bus, IOP and device control unit.

  In order to achieve maximum tape performance, the system must have a balance in the placement of the high performance loads on the system.

  Since the data written to tape must come from the DASD, the DASD and tape operations must be equal or less than the system bus bandwidth. Maximum performance on a 3590 can be gained by having the 3590 on its own bus. Two high performance tape devices on the same bus can have the same effect. During a backup, large LAN or work station networks may also compete with the tape for system bus bandwidth.

  So, if the 3490 or 3590 are placed on a bus with a significant portion of the DASD, both DASD and tape performance will be degraded from possible performance if they are on their own bus.

  Equally important is the tape subsystem bandwidth. If the device has a controller that attaches multiple hosts and multiple devices, they will compete for the device bandwidth in the same manner.

For example, if two 3490 devices are connected to one controller, the controller must share the data path between two devices. This degrades the performance by a factor of two.

- Appending to large capacity tapes

  Large capacity tapes like the 3590 have a very bad effect on the sequential file access. The append process involves locating a spot on the tape to start the write operation. If end positioning of *LEAVE was done on the previous command, saving to the end of tape is not a concern, because the tape is already positioned there. The problem comes when a tape with many files already on it is at the beginning of tape, either from loading a tape or simply rewinding the tape in the drive. In order to position to the end of the tape, OS/400 searches by file marks until it reaches the *END of the tape. The more files on the tape, the longer the search to the end.

  Now there is a new function called Space to End of Data that allows fast access to the end of the tape. With this function, OS/400 can space to the end of a 3590 cartridge (which could hold 30 to 40 gigabytes compressed) in less than 90 seconds.

  To determine if the device supports Space to End of Data function, use the Retrieve Device Capabilities (QTARDCAP) API. See the *System API Programming*, SC41-5800-00 for information about the QTARDCAP API.

  Performance is improved for positioning a tape by specifying a sequence number. This only applies to devices that support the Space to End of data function. A save to a specific sequence number will only be slightly slower than a save to SEQNBR(*END).

  Note: Performance improvements for positioning a tape to a specified sequence number have not been implemented for the CHKTAP, DMPTAP, DSPTAP and DUPTAP commands.

- Excess error recovery

  All tape devices have built in recovery procedures. These recovery procedures degrade performance of the device, when a high number of errors occurs on read and write commands. Extra recovery is very costly to high performance tape devices. As a thorough analysis of a performance problem, a poorly performing device or media should be ruled out by examining the error logs and media statistics. If a tape cartridge has a high number of errors, replace the media. If a tape device has a high number of errors, ensure that the device is cleaned properly.

- DUPTAP performance

  There are many variables to the performance of a Duplicate Tape (DUPTAP) command.

  – Device type

    Using like devices will allow OS/400 to stream devices more easily. Duplicating an 8mm cartridge to a 9348 6250 reel is not a high performing operation. It is very difficult for OS/400 to keep the 9348 streaming. In fact, the 9348 attempts to continue moving and causes a performance degradation known as Hitchback by moving past where the data will be written, having to space backwards, and then write the data. It is very poor use of a tape device to continually force it to stop and start.

  – DUPTAP

    It is slower than doing another save if a save operation is driving a save close to the maximum speed for the device, it is impossible to have a better duplication time than the save because the device is already running as fast as it possibly can.

– Duplicating HDC tapes

Ensure that the tape is duplicated to a device that supports HDC, if the tape has been created with HDC. If the TODEV does not support HDC on DUPTAP and the FROMVOL has been compressed with the DTACPR parameter on a SAVxxx command, then the duplication of the media enforces that a compressed tape be generated. Because compression is created on a device that does not support HDC, SDC is used on the restore of the cartridge and will severely impact performance.

- 1/4-inch device performance

The 1/4-inch tape devices are intended primarily for save and restore operations. These devices are not designed for operations that cannot process data fast enough to keep the tape moving. Performance and reliability will suffer if the tape stops and starts too often. When a tape device is running at the maximum speed of the device and is kept moving, it is known as streaming.

If you use the 1/4-inch cartridge devices for operations other than save and restore, the tape will be more likely to stream if large fixed blocks are used.

- Use large block sizes

Process large tape blocks (BLKLEN parameter). If the records are small, use a blocked format (specify *FB, *DB, *VB, or *VBS on the RCDBLKFMT parameter).

New technology allows some devices to support block sizes greater than the 32K in length. The 3590 devices for instance supports up to 256K on OS/400. Devices also have an optimal block size to be used in conjunction with OS/400 save functions. The USEOPTBLK parameter has been added to the following SAVxxx commands for using this optimal block size:

SAV	SAVDLO
SAVCHGOBJ	SAVCFG
SAVLIB	SAVOBJ
SAVSAVFDTA	SAVSECDTA
SAVSYS	

Specifying USEOPTBLK(*YES) may allow performance gains in your day to day save and restore operations. Specifying USEOPTBLK(*NO) allows you to use a block size close to the 32K block length.

- Use fixed-length records

Use fixed-length records because they are processed more efficiently than variable-length records.

- Design of a high level application for tape

Design the application to do as little processing as possible between tape read or write operations. The best application design is the one that uses a system copy file command (CPYF, CPYFRMTAP, or CPYTOTAP) to copy records between tape and a database file, with application programs that only process records in the database file.

Commonly asked questions about tape usage and support

How do I find out how much data was actually written to a tape with the Save Object (SAVOBJ) command, how to compare compression ratios of a save with or without compression, and how to verify that all data on the tape can be read by the system (at this moment) without a media error?

A tool provided in QUSRTOOL called READTAPE can do all of these.

The following shows a sample output of the tool for a tape with two files on it:

- Data set QGPL at sequence number 1 read.
  - Number of bytes: 25448448
  - Number of blocks: 872
  - Average block length: 29184
- Data set QUSRTOOL at sequence number 2 read.
  - Number of bytes: 13103616
  - Number of blocks: 449
  - Average block length: 29184
- Summary of data read:
  - Total number of data sets: 2
  - Total number of bytes: 38552064
  - Total number of blocks: 1321
  - Average block length: 29184

The following information is provided in QUSRTOOL to explain the READTAPE tool.

The READTAPE tool provides a method for determining the amount of data written on a tape volume.

The READTAPE tool reads all of the data sets on a tape volume. It reports the number of bytes and blocks in each data set. It also computes the average block size which can be used to determine compression ratios on a tape containing compressed save and restore data given that the AS/400 system writes save and restore data in 24,576 byte blocks when saving to a tape device. 24,576 byte blocks are used on the following versions of OS/400 by save and restore: V2R3M0, V3R0M5, V3R1M0, V3R2M0. 29,184 byte blocks are used by V3R6M0, V3R7M0. In V3R7M0, if the device supports an optimum block size, such as the 3590 device, then the optimum value will be used instead of the 29,184 block size. After all of the data sets have been read from the tape volume, overall statistics for the data read are reported to the user.

For example, a 1/2-inch reel tape volume contains 2400 feet of tape. It is initialized to a density of 6250 bits per inch. On it resides a single data set created by the Save Library (SAVLIB) command. The data set is 100 blocks in length for a total of 1105920 bytes of data. The average block length is 11059 bytes, assuming interblock gaps are 1/2 inch in length.

The compression ratio for the data saved is computed by dividing the average block length by the uncompressed block length. 11059/24576 equals 0.45, or 45%. This means the data set used only 45% as much of the tape capacity as it would have used if the data had not been compressed.

To determine the amount of tape used, divide the number of bytes written by the tape density, and add in a factor for the interblock gaps between each block written to the tape. Subtracting this amount from the total length of the tape will give the amount of tape still not used. (1105920/6250)+(0.5*100) equals 227 inches of tape used. The tape is 2400*12 or 28800 inches in length. The amount of tape that is not used is 28800-227 which equals 28573 inches or 2381 feet.

These calculations do not account for temporary errors detected and recovered by the device, and therefore not reported to the system by the tape device while writing data to the tape. They also do not take into consideration tape labels nor

tape marks normally written on a tape volume. The number of temporary errors depends on the cleanliness of the read/write head in the tape device and the quality of the tape media being used. The number of tape labels and tape marks depends on the format of the tape.

Limitations

If an error is detected while the system is reading data from the tape, the tool reports statistics for the data sets that were fully read and the ones that were partially read, and then ends. This tool does not include tape labels processed when computing tape statistics. In typical situations, the amount of tape used by tape labels is insignificant compared to the amount of tape containing data blocks.

Using the Tool

To use the tool, the user must first create the objects that make up the pieces of the tool itself. This is done by creating and executing a CL program that builds the necessary objects. The required source members are located in QUSRTOOL and are listed below.

The following are source members/source file/descriptions of information in QUSRTOOL.
- TTACRT / QATTCL / Source for Install Program
- TTADLT / QATTCL / Source for Delete Program
- TTARTAPE / QATTCMD / Source for READTAPE Command
- TTAREADT / QATTSYSC / Source for READTAPE Program

After the tool is created, the program is started by entering the appropriate command: READTAPE. The following notes describe how to create and use the tool.

Creating the Tool

- Create a library to be used for the installation program. If an existing library is used, this step can be omitted. The characters *crt-lib* are used to represent the installation program library in the following commands.
- Create a library to be used for the tool objects. If the tool objects and the installation program reside in the same library, this step can be omitted. The characters *tool-lib* are used to represent the tool object library in the following commands.
- Create the installation program.
  ```
  CRTCLPGM PGM(crt-lib/TTACRT)
  SRCFILE(QUSRTOOL/QATTCL)
  ```
- Add the tool object library to the library list.
  ```
  ADDLIBLE tool-lib
  ```
- Run the installation program to create the tool objects. CALL crt-lib/TTACRT tool-lib The *tool-lib* must exist and must be a part of the library list.

  If a tool object already exists in the tool library, the tool object will be replaced. The following objects are created in the specified library.
  - READTAPE - *PGM

    C program that reads data from a tape.
  - READTAPE - *CMD

Command that calls READTAPE.

– TTADLT - *PGM

CL Program that deletes the tool objects.

Reading a Tape

To read data from a tape and generate statistics regarding the data, use the READTAPE tool and provide the following parameters:

- DEV - The name of the device on which the tape volume to be read is loaded.
- LABELS - Specifies the type of labeling that is used on the tape volume. The default value *SL is used when a standard labeled tape volume is loaded. The other value which may be specified is *NL and is used when a non-labeled tape volume is loaded.
- ENDOPT - Specifies the end of tape positioning which should take place after the tape volume is processed. The default value *REWIND is used if the tape is rewound. The *UNLOAD value should be used if the tape is unloaded after processing is complete.

Example of Using READTAPE tool

```
READTAPE DEV(TAP01) LABELS(*SL)
ENDOPT(*UNLOAD)
```

When the READTAPE command is run, the statistics regarding the data are written to the standard output (STDOUT) file. If desired, the Override Print File (OVRPRTF) command can be used to redirect the output to a printer file such as QPRINT for later viewing or printing. STDOUT can also be redirected to a physical file if so desired. The READTAPE program creates a tape file object in library QTEMP called TEMP. If a tape file already exists in QTEMP by this name, it will be overlayed with the tape file created by the READTAPE command.

Deleting the Tool

TTADLT is a cleanup program created by the installation program. TTADLT removes objects created by TTACRT (except the library and the program TTADLT) and/or removes the source members used by the READTAPE tool. To run the delete program, enter one of the following.

- `CALL TTADLT (*YES *NONE)`

  If you ONLY want to remove the source members used by the tool.

- `CALL TTADLT (*NO tool-lib)`

  If you ONLY want to remove the objects that the tool created.

- `CALL TTADLT (*YES tool-lib)`

  If you want to remove the source members used by the tool and the objects that the tool created.

  The installation program TTACRT is deleted if it resides in the tool library.

How can I write tape files with a specific creation date?

Because of the meaning of the *creation date*, the system places the current date in the creation date field of the labels when the tape is created. The creation date cannot be overridden by a CL command user. The Duplicate Tape (DUPTAP) command will duplicate a tape exactly, that is, preserving the creation date of the volume being duplicated.

What does the tape or system do if the end of tape marker is crossed while the system is writing trailer labels?

For 1/2-inch technology (technology that allows overwriting of data), the system changes the EOF labels that are written to EOV labels, asks the operator to load the next tape, and writes a null file to the next tape. This null file is written so that the file can be properly extended. For more information on extending data files, see "Extending Files on Tape" on page 9.

Chapter 3. Diskette Support

The program uses Diskettes to create backup copies of information, provide offline storage of files and libraries, or to transfer information to other systems or devices.

The AS/400 system supports both 5 1/4-inch and 8-inch diskettes, with the following restrictions:

- 5 1/4-inch diskettes must be double-sided, high capacity (2HC) diskettes. These diskettes have a capacity equivalent to 8-inch, double density diskettes. When initializing these diskettes, use 2D format and sector sizes 256, 512, or 1024.
- 8-inch diskettes can be any of the following three types:
  - Diskette 1 is a single-sided, single-density diskette.
  - Diskette 2 is a double-sided, single-density diskette.
  - Diskette 2D is a double-sided, double-density diskette.

A double-density diskette contains information on both sides and stores twice the amount of information in the same space as a diskette 1. A diskette 2D holds approximately four times the amount of information as a diskette 1.

For information about how to use the diskette unit for save/restore operations, see the *Backup and Recovery* book.

Related CL Commands

The following commands are available to help maintain and use diskettes. The *CL Reference (Abridged)* manual contains detailed descriptions of these commands.

Diskette Device Description Commands

CHGDEVDKT
Change Device Description (Diskette): The command changes a device description for a diskette unit.

CRTDEVDKT
Create Device Description (Diskette): The command creates a device description for a diskette unit.

DLTDEVD
Delete Device Description: The command deletes a device description.

Diskette Device File Commands

CHGDKTF
Change Diskette File: The command changes certain attributes of a diskette device file.

CRTDKTF
Create Diskette File: The command creates a diskette device file that is used to read and write records on diskette.

DLTF Delete File: The command deletes files.

DSPFD
Display File Description: The command displays the current characteristics of a file.

OVRDKTF

Override with Diskette File: The command overrides a diskette unit file or the file attributes that are specified in a program.

Other Diskette Support Commands

CHKDKT

Check Diskette: This command checks a diskette for a specific volume identifier, file label, or both.

CLRDKT

Clear Diskette: This command clears all files on the diskette by deleting the labels and creates an expired file, labeled DATA. The expired file DATA includes the entire diskette. The system does not delete the data.

CPYFRMDKT

Copy from Diskette: This command copies records from a diskette file to an output file or a printer.

CPYTODKT

Copy to Diskette: The system copies records to a diskette file from a physical, logical, tape, diskette, or spooled inline file.

DLTDKTLBL

Delete Diskette Label: This command deletes a file label from a diskette.

DSPDKT

Display Diskette: This command displays either volume and both file labels or save and restore information on a diskette.

DUPDKT

Duplicate Diskette: The command copies information from one diskette to one or more diskettes.

INZDKT

Initialize Diskette: The command initializes a diskette to write identification information and to format the diskette for use by the system.

RNMDKT

Rename Diskette: The command changes the volume and owner identifiers in the volume label.

Diskette Exchange Types

Diskette data must use one of the following exchange types, and the diskettes must have the attributes (diskette type, sector size, and record length) appropriate for that exchange type. The exchange types are:

* Basic Exchange

  Systems capable of reading and writing both the IBM Diskette 1 and the IBM Diskette 2 can exchange basic exchange data sets. The sector size must be 128 bytes. The length of the records can be from 1 to 128 bytes, with one record per sector.

* H Exchange

  Systems capable of reading and writing the IBM Diskette 2D may exchange type H exchange data sets. The sector size must be 256 bytes. The length of the records can be from 1 to 256 bytes, with one record per sector.

* E Exchange

Type E exchange data sets force the using system to examine each field in the header label. Save/restore operations create E exchange data sets.

- I Exchange

  Type I exchange data sets may be used for:

  - Diskette 1 or Diskette 2 with sector sizes of 128, 256, or 512 bytes
  - Diskette 2D with sector sizes of 256, 512, or 1024 bytes

  The record length can be from 1 to 4096 bytes. One sector may contain one, more than one, or part of a record, or one record may span across sectors, depending on sector size and record length.

  For these exchange types, each record in the file is fixed-length (each record contains the same number of bytes).

Initializing Diskettes

You must initialize and write a volume label on a diskette before you use it on the system. Diskettes are normally initialized and ready for use when you receive them.

The Initialize Diskette (INZDKT) command initializes a diskette to a specified format. You may want to use the INZDKT command to reinitialize a diskette in the following situations:

- If you want to change the sector size on a diskette to a size that is acceptable for a specified exchange type (basic, H, or I).
- To accommodate a save/restore operation.
- If you encounter errors on a diskette while reading or writing to it. Initializing the diskette may correct the errors or assign an alternate cylinder to replace a defective cylinder. A **cylinder** is a set of tracks on a diskette that a read/write head reads without changing position. If you encounter more than two defective cylinders during the process of preparing a diskette, you receive a message that states that the diskette is unusable. If you receive this message, you should discard the diskette.

Initializing a diskette deletes the data that is contained on the diskette. If you want to save the data, you must do so before initializing the diskette. You can use the Duplicate Diskette (DUPDKT) command to copy the entire diskette to another diskette. You use the Copy File (CPYF) command to copy each file on the diskette to a database file or to another media. You may also use the CPYFRMDKT and CPYTODKT commands to copy data from or to a diskette.

The AS/400 system uses data conversion tables derived from the American National Standards Institute, Inc, Document ANSI X3.26-1970.

Multivolume-Diskette Data Files

Depending on the size of your data files, you may place several files on one diskette, called a multifile volume. You call a file a multi-volume diskette data file when a single data file occupies more than one volume.

You must follow these conventions if you use multivolume-diskette data files:

- All volumes of a multivolume set must have the same format (1, 2, or 2D) as the first volume.
- All volumes must have the same sector size (128, 256, 512, or 1024 bytes).
- You must write all volumes in the same character code (EBCDIC or ASCII).

- For an input file, the file exchange types must be the same on each volume of a multivolume set.
- For an input file, the length of the records in the file must be the same on each volume of the file.
- Except for the first volume, you cannot write to multivolume-diskette data files that contain active data files. An **active data file** is a diskette data file that has an expiration date greater than the system date. If you attempt to do this, a message is sent to the system operator that allows the active data files to be ignored (written over).
- You cannot write a data file to a diskette data file that has an extended label area. (The system does not support extended label areas—cylinders that you allocate to contain additional data set labels.)

When you use a diskette for spooled output, make sure that you write the output to the correct diskette. You can do this by specifying the VOL parameter when you spool the output. The advantage of spooling output is that more than one job can be running at the same time while you spool output.

Note: The volume sequence number field in the diskette file label for multivolume-diskette data files has only two positions. 00 indicates Volume 100, after which the numbers return (wrap around) to 01. The program sends an informational message to the system operator each time the numbers return to 01.

To use diskettes for other than saving and restoring, you must observe the following conventions:

- Each volume identifier can be from 1 through 6 alphanumeric characters.
- If you have multivolume-diskette data files, a volume identifier can apply to more than one volume. However, the program does not require you to use the same volume identifier for every volume in a multivolume diskette group.
- You do not write files to diskettes with extended label areas.

Diskette Device Descriptions and Device Files

To access data from a diskette on the system, two objects must exist:

- First, a device description must exist for each diskette unit to describe the device to the system. The Create Device Description (Diskette) (CRTDEVDKT) command creates the description of the diskette unit. The diskette device description contains information such as device address, device name, device type, model number, and features.
- Second, a device file must exist for the diskette unit. The device file describes:
  – How you present input data to a program from the device.
  – How output data is presented to the device from a program.

  You must not confuse diskette unit files with the actual data files on the diskette volumes. The diskette device file provides a link between the application program and the diskette unit for processing the diskettes which contain the data files.

You must program-describe diskette unit files, which means you describe the fields and records in the program that processes the device file. You do not describe the fields and records in the device file itself. It is not necessary to have a separate device file for each diskette unit. You can use a single device file for several different diskette units by using an override command. You can associate any number of diskette unit files with one diskette unit.

Note: You must vary on the device before you use it. This function can be performed automatically when the system is started or you can use the appropriate vary command. See the *Local Device Configuration* book for information about varying on device descriptions.

IBM-Supplied Diskette Device Files

You ship the following diskette unit files with the operating system for your use:

- QDKT (diskette file)
- QDKTSRC (diskette source file)

These files are all program-described data files in library QGPL. The record format names are the same as the file names. The files contain default values for most parameters.

Note: The device (DEV) for these diskette unit files uses the default value *NONE. You must use the CHGDKTF or OVRDKTF commands to specify a diskette unit before using these files.

Example of Creating a Diskette Device File

You can create additional device files to fit your needs. For example, you can create additional device files to accomplish the following:

- Direct output to a special output queue
- Contain the specific volume and label information for a diskette data file that several programs can use.
- Spool output or not spool output

In the following example, you create a diskette device file, DKFILE, to write output to a diskette.

```
CRTDKTF  FILE(DKFILE) DEV(DKT01) CODE(*ASCII)
```

The diskette unit (specified on the DEV parameter) is DKT01, and the system will write the diskette data file in ASCII code. All other parameters on the CRTDKTF command use their default values.

Because you assume the default values in this example, you must:

- Specify the diskette volume on the diskette in another CL command or in each program that uses the device file.
- Specify file label on the diskette in another CL command or in each program that uses the device file.
- Specify the creation date of the data file on the diskette in another CL command or in each program that uses the device file.

Specifying Diskette Device File Parameters

In diskette unit files, you describe, in the application program, the data in each record. The system views each record as one field with a length equal to the record length.

You can specify the following diskette-device-file attributes for the CRTDKTF, CHGDKTF, or OVRDKTF command:

- The following spooling information for output files:

- – Output queue name
- – Maximum number of records to spool for the data file
- – When the system makes output available to a writer program (at job end, at file end, or immediately as it spools)
- – When the system is to hold the output on the output queue until the system operator releases it
- – If the system saves the output after you produce it.

  Note: If you do not spool output, you must ensure that you write output to the device that is assigned to your device file.

- The device to use with device file. You can specify a device name on the create device file command. You can also specify the device name by using the OVRDKTF or CHGDKTF command.
- The wait time for file allocation. The number of seconds the system waits for the file resources to allocate when you open the device file.
- Whether or not you can share the open data path (ODP) for the device file.
- The volume identifiers of the diskettes you use for the device file (that specifies VOL(*NONE) causes you to ignore volume checking).
- The data file label on the diskette.
- The character code (EBCDIC or ASCII) for the data on the diskettes.
- The creation date of an input data file. The program sends a message to the system operator who determines what actions to take if:
  - – The creation date written on the diskette does not match the date specified in the diskette device file.

  You must specify the date in the format that is specified in the system value QDATFMT. However, the system stores the date you specify in the diskette label in the format, yymmdd.
- The exchange type (basic, H, or I) to use when creating an output file on the diskette. You do not use this attribute when processing an input file.
- The expiration date of an output data file on diskette. The expiration date means you cannot write over the data file until the date has expired. The system considers the file protected. You specify the date in the format that is specified by the system value QDATFMT. However, the system stores the specified date in the diskette label as yymmdd.

  If you do not want the data file to expire, specify *PERM on the EXPDATE parameter.

Table 3 on page 47 lists the parameters that apply to diskettes and when you can specify the parameters. The *CL Reference (Abridged)* manual contains detailed information about how to specify these parameters on the CRTDKTF, CHGDKTF, and OVRDKTF commands.

Using Diskette Device Files in High-Level Language Programs

You can access a diskette unit from a program that uses a program-described diskette device file. To use the diskette device file with a program you must:
- Specify the diskette device file name in the program.
- Or use an Override with Diskette File (OVRDKTF) command.

To use the diskette unit directly, you specify the device name and the default SPOOL(*NO) parameter. The high-level language you use determines what diskette parameters you specify in the program. If the parameters are not specified in the file description or program, you can specify them in an OVRDKTF command.

Table 3 lists the parameters that apply to diskettes and where you can specify the parameters.

Table 3. Diskette File Parameters

CL Parameter	Description	Specified on CRTDKTF Command	Specified on OVRDKTF Command	Specified in HLL Programs
FILE	File name	Qualified device file name	Device file name	ILE RPG, COBOL, BASIC, PL/I, and ILE C programming languages
DEV	Device name	*NONE or list of device names	List of device names	
VOL	Volume	*NONE or volume identifier	*NONE or volume identifier	
LABEL	Label	*NONE or file label	File label	PL/I
FILETYPE	File type	*SRC or *DATA		
EXCHTYPE	Exchange type	*STD, *BASIC, *H, or *I	*STD, *BASIC, *H, or *I	
CODE	Code	*EBCDIC or *ASCII	*EBCDIC or *ASCII	COBOL programming language
CRTDATE	Creation date	*NONE or date	*NONE or date	
EXPDATE	Expiration date	*NONE, *PERM, or date	*NONE, *PERM, or date	
SPOOL	Spool data	*NO or *YES	*NO or *YES	
OUTQ	Output queue	Qualified name	Qualified name	
MAXRCDS	Maximum records	*NOMAX or maximum records	*NOMAX or maximum records	
SCHEDULE	Schedule	*FILEEND, *JOBEND or *IMMED	*FILEEND, *JOBEND or *IMMED	
HOLD	Hold	*NO or *YES	*NO or *YES	
SAVE	Save	*NO or *YES	*NO or *YES	
OUTPTY	Output priority	*JOB or output priority	*JOB or output priority	
USRDTA	User data	*BLANK or user data	*BLANK or user data	
IGCDTA	Double-byte data	*NO or *YES	*NO or *YES	
WAITFILE	File wait time	*IMMED, *CLS or number of seconds	*IMMED, *CLS or number of seconds	
SHARE	Shared file	*YES or *NO	*YES or *NO	PL/I
AUT	Authority	*CHANGE, *ALL, *USE, *EXCLUDE, or authorization list name		
REPLACE	Replace existing file	*YES or *NO		
TEXT	Text	*BLANK or text		
	Record length			ILE RPG, COBOL, BASIC, PL/I, and ILE C programming languages

Open Processing for Diskette Device Files

The following considerations apply to opening diskette unit files:

- When you open a diskette unit file, the parameters you specify in the file are merged (overridden) with the parameters you specify in the program. The system merges these parameters with the parameters you specify on an OVRDKTF command, if they are specified.

- You must specify the device name when you open the diskette device file if you do not spool the file; that is, if you specify the default SPOOL(*NO). If you specify DEV(*NONE) in the diskette device file, you must specify the device name on an OVRDKTF command.

- For an input file, the program may specify the record length, but it is not a requirement. If you:
  - do not specify the record length.
  - specify with a length of 0.

  The system will determine the record length from the data file label on the diskette. If the program specifies a record length that is not equal to the length of the records in the data file:
  - The program pads or truncates the records to the length specified in the program.

  The system sends a diagnostic message stating the system pads the records:
  - If the program record length is greater than the length of the records in the data file.

- The program must specify the record length for output files. When the program specifies a record length that exceeds what the diskette format allows:
  - The system truncates the records and sends a diagnostic message to the program.

  Maximum record lengths supported by exchange types are:

  Basic exchange
  > 128 bytes

  H exchange
  > 256 bytes

  I exchange
  > 4096 bytes

- When you open the file the system requires a file label name. If you create the diskette unit file with the LABEL(*NONE) parameter specified, you must supply the label with the OVRDKTF command. The label name must not exceed 8 characters for files in exchange type BASIC, H, or I. The first character must be alphabetic (A through Z, #, $, or @). All remaining characters in the label name can be any character (A through Z, 0 through 9, #, $, or @) except a blank. The period character (.) is valid only in labels on E exchange files. If the period appears in a label on a BASIC, H, or I exchange file, an error is returned.

  The system will not write a data file with a label name that is not valid. If the input file encounters such a label name ,the system sends a diagnostic message and an attempt is made to read the file.

- When the program specifies a source data file type in the diskette device file:
  - The program adds a date and sequence number to each record on input operations.
  - The program removes a date and sequence number from each record on output operations.
  - The date field always remains 0.

The program can specify, if the high-level language you use allows it, for the system to check if the input or output file is a source file. If a program uses a source file that specifies a record length, the file must include 12 bytes for source data. For example, if the data is 80 bytes long, the program must specify a record length of 92.

- A volume identifier specifies the diskette to process. If you do not specify a volume identifier , the system processes the diskette without checking the volume identifier. When you specify a volume identifier, you must load the correct volume. If not, the program sends a message to the system operator that requests you load the correct diskette.

- For multivolume-diskette data files, all volumes must:
  - Have the same type of diskette.
  - Have the same sector size.
  - Have the same record length.
  - Have the same character code.

  For output files, all volumes after the first volume must be written to diskettes that do not contain active files. If active files exist, a message is sent to the system operator that allows the active files to be ignored (overwritten).

- When an output file is created, all expired files on a diskette are written over; no messages are sent to your program. An expired file is a file having an end date less than or equal to the system date.

I/O Processing

The following considerations apply to input/output (I/O) operations performed on diskette data files:

Read and Write Considerations

- When a volume exchange occurs while processing multivolume-diskette data files, a message identifying the current volume identifier is sent to your program.
- If the record length specified in the program does not match the length of the data, the data is padded or truncated to match the record length specified in the program.

Force-End-of-Data Considerations

- The force-end-of-data function is valid for both input and output.
- The force-end-of-data function for an output file does not write any data on diskette.
- The force-end-of-data function for an input file positions the diskette at the last volume of the file and signals end-of-file to the program using the file. Refer to the chapter on spool support in the *Data Management* book for information about how this function is handled for a SPOOL(*YES) output file.

Close Processing

When a diskette device file is closed, the labels for output data are written on the diskette before the file resources are deallocated.

Handling Diskette Errors

If you encounter an error on the input diskette while using the DUPDKT command, the operation is ended without any data being written to the output diskette.

If you encounter an error on the input diskette while using the CPYF, CPYTODKT, or CPYFRMDKT command, all the data in the data file before the record that caused the error is copied. The remaining data in the data file is not copied.

Chapter 4. Overriding Device Files and Device File Attributes

You can use overrides to temporarily change a file name, a device name associated with the file, or some of the other attributes of a file. Overrides allow you to make minor changes in how a program function. Overrides allow you to select the data, on which they operate, without recompiling the program.

Determining Whether or Not to Use Overrides

Use the following properties of overrides to help you decide if overrides are appropriate for the task you want to perform:

- Overrides remain in effect only for the job or program for which they are issued. They do not permanently change the attributes of a file.
- Overrides have no effect on the other jobs that may be running at the same time.
- You must specify and apply overrides before the program openthe file or before you compile a program that opens a file.
- You can enter Override commands interactively from a display station or as part of a batch job.
- you can include Override commands in a control language (CL) program, or you can issue them from other programs via a call to the program QCMDEXC.
- Overrides can be scoped to the call level, activation group level, or the job level.

Tape and Diskette Override Commands

You can process override functions for files using the following CL commands:

OVRDKTF

Override with Diskette File: The Override with Diskette File command does the following:

- It overrides (replaces) the diskette file that is named in the program.
- It overrides certain parameters of a diskette file that is used by the program.
- It overrides the file and certain parameters of the diskette file that is used by the program.

OVRTAPF

Override with Tape File: The Override with Tape File command does the following:

- It overrides (replaces) the tape file that is named in the program.
- It overrides certain attributes of a file that is used by the program.
- It overrides the file and certain attributes of the file.

Overriding File Attributes

File attributes are built as a result of the following:

- Create file command. Initially, this command builds the file attributes.
- Program using the files. At compile time, the user program can specify some of the file attributes. The high-level language used by the program determines the attributes that can be used.

- Override commands. At program run time, these commands can override the file attributes previously built by the merging of the file description and the file parameters specified in the user program.

The simplest form of overriding a file is to override some attribute of the file.

For example, assume that you create a tape file OUTPUT whose attributes are:
- Use Device TAP01.
- Write the Data on the tape with a density of 1600 bits per inch (bpi)
- Use the ASCII character code type.
- When you close the tape file, the program rewinds and unloads the tape.

The Create Tape File (CRTTAPF) command looks like this:

```
CRTTAPF  FILE(QGPL/OUTPUT) DEV(TAP01)
DENSITY(1600) CODE(*ASCII) ENDOPT(*UNLOAD)
```

Your application program has the tape file OUTPUT specified with a character code type of EBCDIC and a density of 3200. However, before you run the application program, you want to change the density to 6250 bpi and the end option to *REWIND. The override command looks like this:

```
OVRTAPF  FILE(OUTPUT) DENSITY(6250)
ENDOPT(*REWIND)
```

When you call the application program, the system uses a tape density of 6250 bpi and the end option is *REWIND.

When the application program opens the file the system merges the following to form the open data path (ODP):
- The file overrides.
- The program-specified attributes.
- The file attributes.

The program uses the open data path (ODP) during the running of the program. File overrides have precedence overprogram-specified attributes. Program-specified attributes have precedence over file-specified attributes. In the example Figure 6 on page 53, when you open the file and you perform output operations, the program writes:
- To the device TAP01 with a density of 6250 bpi.
- In a character code type of EBCDIC.
- An end option of *REWIND.

Figure 6 on page 53 explains this example.
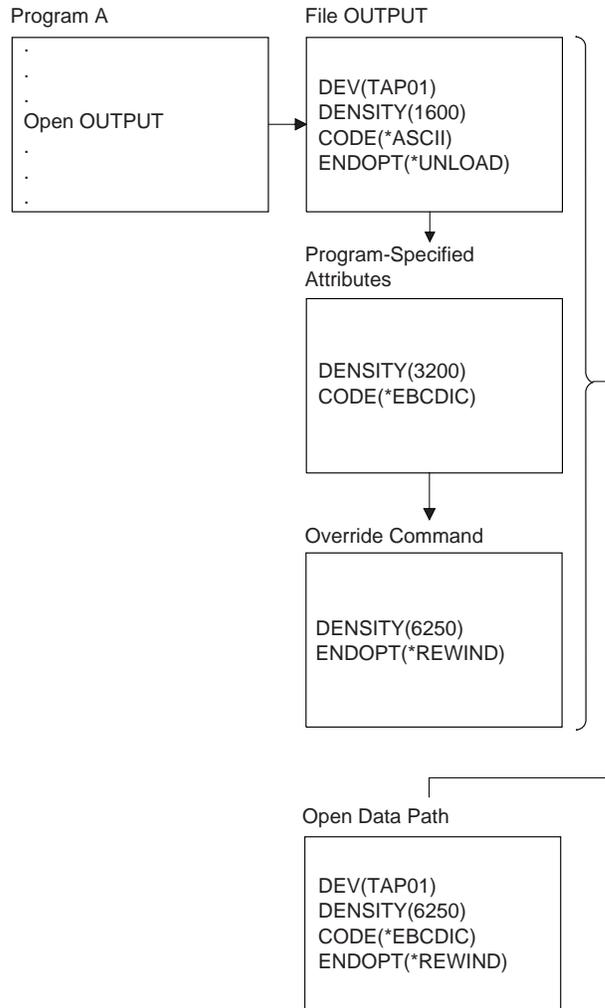
```
   Program A                    File OUTPUT
   ┌─────────────────┐          ┌─────────────────┐
   │ .               │          │                 │
   │ .               │          │ DEV(TAP01)      │
   │ .               │          │ DENSITY(1600)   │
   │ Open OUTPUT     │─────────▶│ CODE(*ASCII)    │
   │ .               │          │ ENDOPT(*UNLOAD) │
   │ .               │          │                 │
   │ .               │          │                 │
   └─────────────────┘          └─────────────────┘
                                         │
                                Program-Specified
                                Attributes
                                ┌─────────────────┐
                                │                 │
                                │                 │
                                │ DENSITY(3200)   │
                                │ CODE(*EBCDIC)   │
                                │                 │
                                │                 │
                                │                 │
                                └─────────────────┘
                                         │
                                Override Command
                                ┌─────────────────┐
                                │                 │
                                │                 │
                                │ DENSITY(6250)   │
                                │ ENDOPT(*REWIND) │
                                │                 │
                                │                 │
                                └─────────────────┘

                                Open Data Path
                                ┌─────────────────┐
                                │                 │
                                │ DEV(TAP01)      │
                                │ DENSITY(6250)   │
                                │ CODE(*EBCDIC)   │
                                │ ENDOPT(*REWIND) │
                                │                 │
                                └─────────────────┘

                                          RV2W702-1
```

Figure 6. Overriding File Attributes

Overriding File Names in HLL Programs

Another simple form of overriding a file is to change the file that is used by the
program. This can be useful for moved or renamed files after you compile the
program. For example, you want to send the output from your application program
to tape file TAPE20 instead of the tape file OUTPUT1 (The application program
specifies OUTPUT1). Before you run the program, enter the following:

```
OVRTAPF FILE(OUTPUT1) TOFILE(TAPE20)
LABEL(FILE01)
```

You must create the file TAPE20 by a CRTTAPF command before you use it.

Figure 7 on page 54 explains this example.
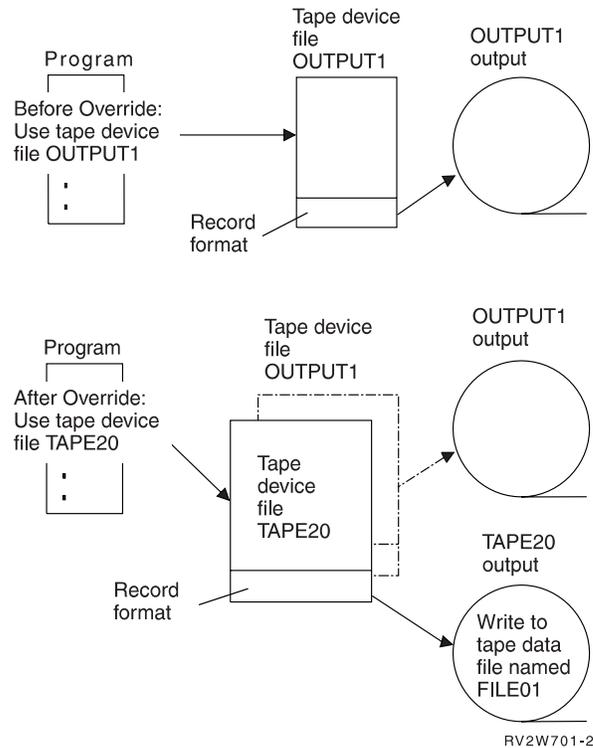
Figure 7. Overriding a File Name

You may want to override a file with a file that has a different file type. For example, you may want to override a diskette file with a display file. To determine if you can override your file with a file that has a different type, see the information under "File Redirection" on page 55. More information is available in the *Data Management* book.

Overriding Both File Names and File Attributes in HLL Programs

This form of overriding files is simply a combination of overriding file attributes and overriding file names or types. This form of override can override the file used in a program. You can also override the attributes of the overriding file.

For example, you want the output from your application program directed to the tape file REPORTS instead of the tape file OUTPUT1 (Your application program specifies OUTPUT1). In addition to having the application program use the tape file REPORTS, you wish to override the density to 6250 bpi. Assume you create the file REPORTS with the following command:

```
CRTTAPF FILE(REPORTS) DEV(TAP01) DENSITY(1600)
```

Before you run the program, type the following command:

```
OVRTAPF FILE(OUTPUT1) TOFILE(REPORTS)
DENSITY(6250)
```

When you call the application program you write the data using the tape file REPORTS with a density of 6250 bpi.

This is *not* equal to the following two override commands:

Override 1
```
OVRTAPF FILE(OUTPUT1) TOFILE(REPORTS)
```

Override 2
```
OVRTAPF FILE(REPORTS) DENSITY(6250)
```

The program applies only one override for each call level for an open of a particular file. You use a single command to override the file used by the program and override the attributes of the overriding file from one call level. If you use two overrides, override 1 causes you to direct the output to the tape file REPORTS, but you ignore override 2.

Deleting Overrides

If you want to delete an override, you can use the Delete Override (DLTOVR) command.

The DLTOVR command, in an application that either calls or transfers control to other programs, may or may not delete the override. More information about deleting overrides in application programs is available in the *Data Management* book.

Displaying Overrides

You can display all file overrides or file overrides for a specific file using the Display Override (DSPOVR) Command.

The DSPOVR command displays the overrides used by an application that either calls or transfers control to other programs. You can control which overrides are displayed. More information about displaying overrides used in application programs is available in the *Data Management* book.

File Redirection

File redirection refers to using overrides to change the file name and library or the type of the file you process. For example, you can substitute:
- One tape file for another tape file.
- One diskette file for another diskette file.
- Or change from using a tape or diskette file to using a display file, printer file, ICF file and so on.

System code may or may not support file redirection. See "Recognizing Commands that Ignore or Restrict Overrides" on page 58 for rules on how system code processes overrides.

Overriding Files with the Same File Types

When a program replaces a file that is used with another file of the same type:
- The system processes the new file in the same manner as the original file.

If you redirect a field level file, or any other file containing externally described data:
- It is usually necessary to specify the command LVLCHK(*NO) or recompile the program.

With level checking that is turned off, it is still necessary that the record formats in the file be compatible with the records in the program. If the formats are not compatible, you cannot predict the results.

Overriding Files with Different File Types

If you change to a different type of file, the system ignores the device-dependent characteristics and reads or writes the records sequentially. The program must specify some device parameters in the new device file or the override. The system takes default parameters for others. This manual describes the effect of specific redirection combinations later in this section.

The system ignores any attributes specified on overrides of a different file type than the final file type. The parameters SPOOL, SHARE, OPNSCOPE, and SECURE are exceptions to this rule. The system accepts the above parameters from any override to the file, regardless of device type.

Some redirection combinations present special problems due to the specific characteristics of the device. In particular:

- Do not use file redirection for save files.
- You can redirect nonsequentially processed database files only to another database file or a DDM file.
- You can redirect display files and ICF files that use multiple devices (MAXDEV or MAXPGMDEV > 1) only to a display file or ICF file.
- Redirecting a display file to any other file type, or another file type to a display file, requires:
  - That the program be recompiled, with the override active, if there are any input-only or output only fields.

  This is necessary because the display file omits these unused fields from the record buffer, but other file types do not.

Table 4 on page 56 summarizes valid file redirections.

Table 4. File Redirections

To-File	From-File					
	Printer	**ICF**	**Diskette**	**Display**	**Database**	**Tape**
Printer	O*	O	O	O	O	O
ICF		I/O		I/O		
	O	O	O	O	O	O
		I	I	I	I	I
Diskette	O	O	O	O	O	O
		I	I	I	I	I
Display		I/O		I/O		
	O	O	O	O	O	O
		I	I	I	I	I
Database	O	O	O	O	O	O
		I	I	I	I	I

To-File	From-File					
	Printer	**ICF**	**Diskette**	**Display**	**Database**	**Tape**
Tape	O	O	O	O	O	O
		I	I	I	I	I

- I=input file O=output file I/O=input/output file
- *=redirection to a different type of printer

identify the file type to override in the FROM-FILE columns and the file type to override in the TO-FILE column. The intersection specifies an I or O or both, meaning that the substitution is valid for these two file types when used as input files or as output files.

For instance, you can override a diskette output file with a tape output file, and a diskette input file with a tape input file. The chart refers to file type substitutions only. That is, you cannot change the program function by overriding an input file with an output file.

The following charts describe the specific defaults taken and what to ignore for each redirection combination.

From: Diskette Input

To: ICF: Records are retrieved from the ICF file one at a time.

Display: Records are retrieved from the display one at a time. Type in the data for each record and press the Enter key when the record is complete. A nonfield-level device file must be specified. Diskette label information is ignored.

Database: Records are retrieved in sequential order. Diskette label information is ignored.

Tape: Records are retrieved in sequential order. If a label value is specified in the program, that value is used as the label for the tape file.

From: Diskette output

To: ICF: Records are written to the ICF file one at a time.

Database: Records are written to the database in sequential order.

Display: Records are written to the display with each record overlaying the previous record. You can request each output record using the Enter key.

Printer: Records are printed and folding or truncating is performed as specified in the printer file.

Tape: Records are written on tape in sequential order.

From: Tape input

To: ICF: Records are retrieved from the ICF file one at a time.

Display: Records are retrieved from the display one at a time. Type in the data for each record and press the Enter key when the record is complete. A nonfield-level device file must be specified. Tape label information is ignored.

Database: Records are retrieved in sequential order. One record is read as a single field. Tape label information is ignored.

Diskette: Records are retrieved in sequential order. If a label value is specified in the program, that value is used as the label for the diskette file.

From: Tape Output

To: Printer: Records are printed, and folding or truncating is performed as specified in the printer file.

ICF: Records are written to the ICF file one at a time. Tape label information is ignored.

Diskette: The amount of data written on diskette depends on the exchange type of the diskette. If a label value is specified in the program, that value is used as the label for the diskette file. Refer to Chapter 3. Diskette Support for a description of exchange types.

Display: Records are written to the display with each record overlaying the previous record. You can request each output record using the Enter key.

Database: Records are written to the database in sequential order.

Recognizing Commands that Ignore or Restrict Overrides

The following commonly used commands ignore overrides entirely:

ALCOBJ	DSPFD
CHGOBJOWN	DSPFFD
CPYIGCTBL	EDTOBJAUT
CRTDKTF	GRTOBJAUT
CRTDUPOBJ	MOVOBJ
CRTTAPF	RNMOBJ
DLCOBJ	RVKOBJAUT
DLTF	

The system does not apply overrides to any system files that you open as part of an end-of-routing step or end-of-job processing. For example, you cannot specify overrides for the job log file. In some cases, when you need to override something in a system file, you may be able to change it through a command other than an override command. For example, to change the output queue for a job log, the output queue could be changed before sign-off using the OUTQ parameter on the Change Job (CHGJOB) command to specify the name of the output queue for the job. If the printer file for the job log contains the value *JOB for the output queue, the output queue is the one specified for the job. The save and restore commands open the tape file with SECURE(*YES) so that the tape file overrides will be ignored.

The following commands allow overrides for the SRCFILE and SRCMBR parameters only:

CRTCMD	CRTPF
CRTICFF	CRTPRTF
CRTDSPF	CRTSRCPF
CRTLF	CRTTBL
	CRTXXXPGM

(All create program commands. These commands also use overrides to determine which file will be opened by a compiled program. See *Data Management* book for more information about applying overrides when compiling a program.)

The following commands allow overrides, but do not allow changing the MBR to *ALL:

CPYFRMPCD   CPYTOPCD

The following commands do not allow overrides to be applied to the display files they use. Overrides to the printer files they use should not change the file type or the file name. Various restrictions are placed on changes that may be made to printer files used by these commands, but the system cannot guarantee that all combinations of possible specifications will produce an acceptable report.

DMPTAP
(In addition to the preceding limitations, this command does not allow overrides to the file it dumps.)

DSPDKT and DSPTAP
(The display commands that display information about a file do not allow overrides to that file.)

Appendix. Feedback Area Layouts

This appendix contains Product-Sensitive Programming Interface and Associated Guidance Information.

Figures in this appendix describe the open and I/O feedback areas associated with tape or diskette files. The program presents the following information for each item in these feedback areas:

- Offset, which is the number of bytes from the start of the feedback area to the location of each item
- Data Type
- The program gives the length in number of bytes
- Contents, which is the description of the item and the valid values for it
- File type, which is an indication of what file types are valid for each item.

For more information about feedback area layouts for files other than tape and diskette files, refer to the *Data Management* book.

The support provided by the high-level language you are using determines how to access this information and how the program represents the data types. See your high-level language manual for more information.

Open Feedback Area

The open feedback area is the part of the open data path (ODP) that contains general information about the file after the program opens the file. It also contains file-specific information, depending on the file type. The open feedback area contains information about each device or communications session defined for the file. The program sets this information during open processing and the program may update the information as it performs other operations.

Table 5. Open Feedback Area

Offset	Data Type	Length	Contents	File Type
0	Character	2	Open data path (ODP) type:	Tape and diskette
			DS Display, tape, ICF, save, printer or diskette file that is not being spooled.	
			DB Database member.	
			SP Printer or diskette file being spooled or inline data file.	
2	Character	10	Name of the file being opened. If the ODP type is DS, this is the name of the device file or save file. If the ODP type is SP, this is the name of the device file or the inline data file.	Tape and diskette
12	Character	10	Name of the library containing the file. For an inline data file, the value is *N.	Tape and diskette
22	Character	10	Name of the spooled file. The name of a database file containing the spooled input or output records.	Diskette being spooled
32	Character	10	Name of the library in which the spooled file is located.	Diskette being spooled
42	Binary	2	Spooled file number.	Diskette being spooled
44	Binary	2	Maximum record length.	Tape and diskette

Table 5. Open Feedback Area (continued)

Offset	Data Type	Length	Contents	File Type
46	Character	2	Reserved.	
48	Character	10	Member name.	Diskette
			If ODP type SP, the member name in the file named at offset 22.	
58	Binary	4	Reserved.	
62	Binary	4	Reserved.	
66	Binary	2	File type:	Tape and diskette
			1 Display	
			2 Printer	
			4 Diskette	
			5 Tape	
			9 Save	
			10 DDM	
			11 ICF	
			20 Inline data	
			21 Database	
68	Character	3	Reserved.	
71	Binary	2	Not applicable to tape and diskette.	
73	Binary	2	Not applicable to tape and diskette.	
75	Binary	4	Not applicable to tape and diskette.	
79	Character	2	Not applicable to tape and diskette.	
81	Character	1	Not applicable to tape and diskette.	
82	Character	1	Source file indication.	Tape and diskette
			Y File is a source file.	
			N File is not a source file.	
83	Character	10	Reserved.	
93	Character	10	Reserved.	
103	Binary	2	Offset to volume label fields of open feedback area.	Tape and diskette
105	Binary	2	Maximum number of records that can be read or written in a block when using blocked record I/O.	Tape and diskette
107	Binary	2	Not applicable to tape and diskette.	
109	Binary	2	Blocked record I/O record increment. Number of bytes that must be added to the start of each record in a block to address the next record in the block.	Tape and diskette
111	Binary	4	Reserved.	
115	Character	1	Miscellaneous flags.	
			Bit 1: Reserved.	
			Bit 2: File shareable	Tape and diskette
			0 File was not opened shareable.	
			1 File was opened shareable (SHARE(*YES)).	

Table 5. Open Feedback Area (continued)

Offset	Data Type	Length	Contents	File Type
			Bit 3: Not applicable to tape and diskette.	
			Bit 4: Not applicable to tape and diskette.	
			Bit 5: Not applicable to tape and diskette.	
			Bit 6: Field-level descriptions This is always 0 for tape and diskette.	Tape and diskette
			Bit 7: DBCS-capable file **0** File is not DBCS-capable. **1** File is DBCS-capable.	Tape and diskette
			Bit 8: Not applicable to tape and diskette.	
116	Character	10	Not applicable to tape and diskette.	
126	Binary	2	File open count. If the file has not been opened shareable, this field contains a 1. If the file has been opened shareable, this field contains the number of programs currently attached to this file.	Tape and diskette
128	Binary	2	Reserved.	
130	Binary	2	Not applicable to tape and diskette.	
132	Character	1	Miscellaneous flags.	
			Bit 1: Not applicable to tape and diskette	
			Bit 2: Not applicable to tape and diskette.	
			Bit 3: Not applicable to tape and diskette.	
			Bit 4: Not applicable to tape and diskette.	
			Bit 5: Not applicable to tape and diskette.	
			Bit 6: User buffers **0** System creates I/O buffers for the program. **1** User program supplies I/O buffers.	Tape and diskette
			Bits 7: Reserved.	
			Bits 8: Not applicable to tape and diskette.	

Table 5. Open Feedback Area (continued)

Offset	Data Type	Length	Contents	File Type
133	Character	2	Open identifier. The value is unique for a full open of a file (SHARE(*NO)) or the first open of a file with (SHARE(*YES)). This is used for display and ICF files, but is set up for all file types. It allows you to match this file to an entry on the associated data queue.	Tape and diskette
135	Binary	2	Maximum record format length, including both data and file-specific information such as: first-character forms control, option indicators, response indicators, source sequence numbers, and program-to-system data. If the value is zero, then use the field at offset 44.	Tape and diskette
137	Character	2	Not applicable to tape and diskette.	
139	Character	1	Not applicable to tape and diskette.	
140	Character	6	Reserved.	
146	Binary	2	Number of devices defined for this ODP. For displays, this is determined by the number of devices defined on the DEV parameter of the Create Display File (CRTDSPF) command. For ICF, this is determined by the number of program devices defined or acquired with the Add ICF Device Entry (ADDICFDEVE) or the Override ICF Device Entry (OVRICFDEVE) command. For all other files, it has the value of 1.	Tape and diskette
148	Character		Device name definition list. See "Device Definition List" for a description of this array.	Tape and diskette

Device Definition List

The device definition list part of the open feedback area is an array structure. Each entry in the array contains information about each device or communications session that you attach to the file. The number at offset 146 of the open feedback area shows the number of entries in this array. The device definition list begins at offset 148 of the open feedback area. The offsets shown for it are from the start of the device definition list rather than the start of the open feedback area.

Table 6. Device Definition List

Offset	Data Type	Length	Contents	File Type
0	Character	10	Program device name. For database files, the value is DATABASE. For printer or diskette files being spooled, the value is *N. For save files, the value is *NONE. For ICF files, the value is the name of the program device from the ADDICFDEVE or OVRICFDEVE command. For all other files, the value is the name of the device description.	Tape and diskette
10	Character	50	Reserved.	
60	Character	10	Device description name. For printer or diskette files being spooled, the value is *N. For save files, the value is *NONE. For all other files, the value is the name of the device description.	Tape and diskette

Table 6. Device Definition List (continued)

Offset	Data Type	Length	Contents	File Type
70	Character	1	Device class.	Tape and diskette
			hex 01 Display	
			hex 02 Printer	
			hex 04 Diskette	
			hex 05 Tape	
			hex 09 Save	
			hex 0B ICF	
71	Character	1	Device type.	
			hex 08 Spooled	
			hex 1A 9347 Tape Unit	
			hex 1B 9348 Tape Unit	
			hex 1C 9331-1 Diskette Unit	
			hex 1D 9331-2 Diskette Unit	
			hex 2A 6346 Tape Unit	
			hex 2B 2440 Tape Unit	
			hex 2C 9346 Tape Unit	
			hex 2D 6331 Diskette Unit	
			hex 2E 6332 Diskette Unit	
			hex 3A 3430 Tape Unit	
			hex 3B 3422 Tape Unit	
			hex 3C 3480 Tape Unit	
			hex 3D 3490 Tape Unit	
			hex 49 6367 Tape Unit	
			hex 4A 6347 Tape Unit	
			hex 4E 6341 Tape Unit	
			hex 4F 6342 Tape Unit	
			hex 50 6133 Diskette Unit	
			hex 53 6366 Tape Unit	
			hex 54 7208 Tape Unit	
			hex 5A 6343 Tape Unit	
			hex 5B 6348 Tape Unit	
			hex 5C 6368 Tape Unit	
			hex 64 6344 Tape Unit	
			hex 65 6349 Tape Unit	
			hex 66 6369 Tape Unit	
			hex 67 6380 Tape Unit	
			hex 68 6378 Tape Unit	
			hex 69 6390 Tape Unit	
			hex 70 6379 Tape Unit	
			hex 71 9331-11 Tape Unit	
			hex 72 9331-12 Tape Unit	
			hex 73 3570 Tape Unit	
			hex 74 3590 Tape Unit	
			hex 75 6335 Tape Unit	
			hex 76 1/4-inch Cartridge Tape[1]	
			hex 77 1/2-inch Cartridge Tape[1]	
			hex 78 1/2-inch Reel Tape[1]	

Table 6. Device Definition List (continued)

Offset	Data Type	Length	Contents	File Type
			hex 79 8mm Cartridge Tape[1]	
72	Binary	2	Not applicable to tape and diskette.	
74	Binary	2	Not applicable to tape and diskette.	
76	Character	2	Not applicable to tape and diskette.	
78	Character	1	Not applicable to tape and diskette.	
79	Character	1	Not applicable to tape and diskette.	
80	Character	50	Reserved.	

Note:

1. If the device is not listed above, use one of the following general categories:
   - hex 76—1/4″ Cartridge Tape
   - hex 77—1/2″ Cartridge Tape
   - hex 78—1/2″ Reel Tape
   - hex 79—8mm Cartridge Tape

Volume Label Fields

Table 7. Volume Label Fields

Offset	Data Type	Length	Contents	File Type
0	Character	128	Volume label of current volume.	Tape and diskette
128	Character	128	Header label 1 of the opened file.	Tape and diskette
256	Character	128	Header label 2 of the opened file.	Tape

IBM Standard Volume Label (VOL1)

Table 8. Format of the IBM Standard Volume Label (VOL1) for Tape

Offset	Data Type	Length	Contents	File Type
0	Character	3	Label identifier	
3	Character	1	Label number	
4	Character	6	Volume identifier (Volume Serial Number)	
10	Character	1	Volume access (security)	
11	Character	5	VTOC Pointer (not used)	
16	Character	21	Reserved	
37	Character	14	Owner Name and Address Code (Owner identifier)	
51	Character	29	Reserved	

The IBM standard volume label (VOL1) is 80 characters in length. The system uses it to identify the tape volume, tape volume owner, and security of the tape volume's contents. It is always the first block of data on the tape volume if the tape is a standard labeled tape. The program records the volume label in EBCDIC.

The contents and function of each of the fields is described below.

- Label identifier

   The characters VOL identify this label as a volume label. The system reads this field to verify that it mounts a standard labeled tape. The system also verifies that this label is a volume label. The system writes this field to the tape when you use the Initialize Tape (INZTAP) command and specify the new volume (NEWVOL) parameter.

- Label number

   The relative position of the label within a set of labels of the same type. The label number is always 1 for the IBM standard volume label.

- Volume identifier (Volume Serial Number)

  A unique identification code to identify the logical tape volume. The system writes the value that is specified for the new volume (NEWVOL) parameter when you use the Initialize Tape (INZTAP) command. For media library devices, the program recommends that you match the logical volume identifier with the external bar code identifier on the cartridges. The value may range from 1 to 6 alphanumeric characters (left justified and padded with blanks if less than 6). The alphanumeric character set includes A-Z, 0-9, @, $, and #. If a program specifies a value for a command in the VOL parameter, the system verifies this field to match the specified value.

- Volume access (security)

  The program considers the volume secure (from processing) if the volume security field is a blank, character zero, or hex zero. *SECOFR authority can process the secured volumes.

- VTOC Pointer (not used)

  Not used by OS/400.

- Owner Name and Address Code (Owner identifier)

  The owner identifier of the tape volume. The system writes a value to this field when you use the OWNER parameter on the Initialize Tape (INZTAP) command. The purpose of the field is to write the owner identifier of the volume or to write information about the contents of the volume.

IBM Standard Data Set Label 1 (HDR1/EOV1/EOF1)

Table 9. Format of the IBM Standard Data Set Label 1 (HDR1/EOV1/EOF1)

Offset	Data Type	Length	Contents	File Type
0	Character	3	Label identifier	
3	Character	1	Label number	
4	Character	17	Data Set (File) Identifier	
21	Character	6	Aggregate volume identifier	
27	Character	4	Aggregate volume sequence number	
31	Character	4	Data Set (File) sequence number	
35	Character	4	Generation Number (not used)	
39	Character	2	Generation Version Number (not used)	
41	Character	6	Creation Date	
47	Character	6	Expiration Date	
53	Character	1	Data Set (File) Security (not used)	
54	Character	6	Block Count, Low Order (Trailer labels only)	
60	Character	13	System Code (Trailer labels only)	
73	Character	3	Reserved	
76	Character	4	Block Count, High Order (Trailer labels only)	

The IBM standard data set label 1 (HDR1/EOV1/TRL1) is 80 characters in length, and you use it

To identify each data set. The program records the data set label in EBCDIC.

The program describes the contents and function of each of the fields below.

- Label identifier

  The characters identify the type of data set label.
  - HDR - Header Label (the beginning of a data set)
  - EOF - Trailer Label (the end of a data set)

– EOV - Trailer Label (the end of a data set that is continued on another volume)

- Label number

  The relative position of this label, within a set of labels of the same type; it is always 1 for the IBM data set label 1.

- Data Set Identifier (filename)

  A unique identification code to identify the data set (file). If the data set ID is less than 17 characters, it is left justified, and you pad it with blanks.

- Aggregate volume identifier

  This field contains the volume identifier from the volume labels of the first volume in a multivolume data set.

- Aggregate volume sequence number

  This field contains the relative volume number of this volume in a multivolume data set. The field is 0001 for a single volume data set.

- Data Set (File) sequence number

  This field indicates the relative position of the data set within a multiple data set group. The value will be in EBCDIC displayable characters for 0001–9999. For numbers larger than 9999 that will not fit in the 4–character field as an EBCDIC displayable character set, the first byte will a '?' ('6F'x) for EBCDIC labels. The last three bytes will be a binary number from 1 to 64000.

- Generation Number that is not used.

  If the data set is a part of a generation data group, this field contains a number that indicates absolute generation number. The OS/400 does not use this field.

- Generation Version Number (not used)

  If the data set is part of a generation data group, this field contains a number that indicates the version number of the generation. OS/400 does not use this field.

- Creation Date

  The creation date of the data set. The program shows the date in the format cyyddd, where:

  c = century (blank=19; 0=20; 1=21; and so on)

  yy = year (00-99)

  ddd = day of the current year (001-366)

  Note that the century code gives the first two digits of the year, not the actual century. For example, a blank which translates into 19 indicates a year in the 1900s, not in the nineteenth century.

- Expiration Date

  The date the program considers the data set expired. Expiration date refers to a data set that is allowed to be overwritten. The user specifies the date in the open tape file that is used in writing the tape. You ignore the expiration date on input, and verify the expiration date on output by OS/400. The date is in the format cyyddd, where:

  c = century (blank=19; 0=20; 1=21; and so on)

  yy = year (00-99)

  ddd = day of the current year (001-366)

  Note that the century code gives the first two digits of the year, not the actual century. For example, a blank which translates into 19 indicates a year in the 1900s, not in the nineteenth century.

- Data Set (File) Security that is not used.

A code number indicates the security status of the data set. OS/400 does not use this field. The following values indicate data set security values that are created by other systems:

0   No password protection

1   Password protection (required for read/write/deletion)

3   Password protection that is required for write/deletion.

- Block Count, Low Order (Trailer labels only)

  The field in the trailer labels contains the low order six digits of the number of data blocks in the data set on the current volume. The field in the header labels contains hex zeros.

- System Code (Trailer labels only)

  A unique code that identifies the system that created the data set.

  IBMOS400
        IBM OS/400

  IBM OS/VS 370
        IBM MVS operating system

- Block Count, High Order (Trailer labels only)

  The field in the trailer labels contains the high order four digits of the number of data blocks in the data set on the current volume. The field in the header labels contains hex zeros.

IBM Standard Data Set Label 2 (HDR2/EOV2/EOF2)

Table 10. Format of the IBM Standard Data Set Label 2 (HDR2/EOV2/EOF2)

Offset	Data Type	Length	Contents
0	Character	3	Label identifier
3	Character	1	Label number
4	Character	1	Record Format
5	Character	5	Block Length
10	Character	5	Record Length
15	Character	1	Tape Density/Format
16	Character	1	Data Set Position
17	Character	17	Job/Job Step Identification
34	Character	2	Tape Recording Technique
36	Character	1	Control Character
37	Character	1	Reserved
38	Character	1	Block Attribute
39	Character	3	Reserved
42	Character	5	Device Serial Number (not used)
47	Character	1	Checkpoint Data Set Identifier
48	Character	22	Reserved
70	Character	10	Large Block Length

The IBM standard data set label 2 (HDR2/EOV2/TRL2) is 80 characters in length and the program uses it to identify additional information about the data set. The program records the data set label in EBCDIC.

The program describes the contents and function of each of the fields below.

- Label identifier

  The characters identify the type of data set label.

  HDR   Header Label (the beginning of a data set)

EOF   Trailer Label (the end of a data set)

EOV   Trailer Label (the end of a data set that is continued on another volume)

- Label number

  The relative positions of this label within the set of labels of the same type; The Label number is always 2 for the IBM data set label 2.

- Record Format

  An alphabetic character that indicates the format of the records in the data set. While the operating system reads from the tape, the Record format field tells the operating system how to interpret the blocks of data the program reads.

  F   Fixed length records

  V   Variable length records

  U   Undefined length records

- Block Length

  A number indicating the block length (in bytes) of the data blocks on the tape. The number in this field can range from 18 to 32767 on OS/400. For numbers greater than 32767, the Large Block Length field allows values up to 512 Kbytes.

- Record Length

  A number that indicates the record length, in bytes, of the logical records on the tape volume. The interpretation of the number depends on the Record Format field.

  F   Fixed length records.

  V   Variable length records.

  U   Undefined length records

- Tape Density/Format

  A code indicating the record density/format of the tape volume.

  3   1600 bpi

  4   6250 bpi

  5   3200 bpi

  blank   all other densities/formats

- Data Set Position

  A code indicating a volume switch is as follows:

  0   No volume switch has occurred

  1   A volume switch previously occurred

- Job/Job Step Identification

  This field identifies the job and job step that created or extended the data set. OS/400 does not use this field.

- Tape Recording Technique

  This field indicates the tape recording technique used in creating the data set.

  blank   No Improved Data Recording Capability (IDRC) used.

  'P '   Improved Data Recording Capability (IDRC) used.

- Control Character

  A printer control code indicating whether the program uses a control character set to create the data set and the type of control characters used:

A   Contains ANSI control characters

M   Contains machine control characters

blank   Contains no control characters

- Block Attribute

  A code indicating the block attribute used to create the data set:

  B   Blocked records

  S   Spanned records, if the record format byte is V

  S   Standard records, if the record format byte is F

  R   Blocked and spanned records, if the record format byte is V

  R   Blocked and standard records, if the record format byte is F

  blank   Records that are not blocked and not spanned, or records that are not blocked and not standard

- Device Serial Number (not used)

  The serial number of the device that writes the volume. Header and trailer labels may have different serial numbers if the program extends the data set. OS/400 does not use this field.

- Checkpoint Data Set Identifier

  This byte contains the character C if the data set is a secure checkpoint data set. The byte is blank if the data set is not a secure data set checkpoint.

- Large Block Length

  A number indicating the block length (in bytes) of the data blocks on the tape. The number in this field can range from 18 to 524288 on OS/400. For numbers up to 32767, the Block Length field also contains the block length in bytes.

IBM Standard User Labels (UHL1-UHL8 or UTL1-UTL8)

Table 11. Format of the IBM Standard User Labels (UHL1-UHL8 or UTL1-UTL8)

Offset	Data Type	Length	Contents
0	Character	3	Label identifier
3	Character	1	Label number
4	Character	76	User data

The IBM standard user labels (UHL1-UHL8 or UTL1-UTL8) is 80 bytes in length. The program uses them to identify additional information about the data set. All information in the user labels is user data. The user labels directly follow the header/trailer group and can be up to 8 user labels per data set.

The program describes the contents and function of each of the fields below.

- Label identifier

  The characters identify the type of data set label.

  UHL   User Header Label (the beginning of a data set)

  UTL   User Trailer Label (the end of a data set)

- Label number

  The relative position of this label within the set of labels of the same type; there is a limit of 8 user labels per data set.

- User Data

  You contain any user information in the user label. The program ignores the information in this set by OS/400, but passes it to a specified user label program.

Other IBM Standard Labels

OS/400 only supports labels described above. However, it reads and ignores volume labels (VOL2-VOL8) and file labels (HDR3-HDR8) created on other operating systems.

ISO/ANSI Standard Volume Label (VOL1)

Table 12. Format of the ISO/ANSI Standard Volume Label (VOL1), Version 3

Offset	Data Type	Length	Contents	File Type
0	Character	3	Label identifier	
3	Character	1	Label number	
4	Character	6	Volume identifier (Volume Serial Number)	
10	Character	1	Accessibility	
11	Character	26	Reserved	
37	Character	14	Owner Identifier	
51	Character	28	Reserved	
79	Character	1	Label Standard Level	

The ISO/ANSI standard volume label (VOL1) is 80 characters in length. The program uses the label to identify the tape volume, tape volume owner, and security of the tape volume's contents. It is always the first block of data on the tape volume if the tape is a standard labeled tape. The field definition follows the industry standards as understood and interpreted by IBM: (OS/400 supports for input only):

- ANSI X3.27-1978, level 4
- ISO 1001-1979,level 4

The system records the volume label in ASCII.

the program describes the contents and function of each of the fields below. This version of the ISO/ANSI standard is Version 3.

- Label identifier

  The characters VOL identify this label as a volume label. The system reads this field to verify that a standard labeled tape is mounted, and that this label is a volume label. When you use the initialized tape (INZTAP) command and specify the new volume (NEWVOL) parameter, the system writes this field to the tape.

- Label number

  The relative position of this label within the set of labels of the same type. The Label number is always 1 for the Version 3 volume label.

- Volume identifier

  A unique identification code to identify the logical tape volume. The system writes the value specified for the new volume (NEWVOL) parameter on the Initialize Tape (INZTAP) command in this position. For media library devices, the logical volume identifier should match the external bar code identifier on the cartridges. The value may be from 1 to 6 alphanumeric characters (left justified and padded with blanks if less than 6). The alphanumeric character set includes A-Z, 0-9, @, $, and #. If the program specifies a value for a command in the VOL parameter, the system verifies that this field matches the specified value.

- Accessibility

  The system considers the volume secure (from processing) if the volume security field is a blank. A program with *SECOFR authority can process secured volumes.

- Owner Identifier

The owner identifier of the tape volume. The system writes a value to this field through the OWNER parameter on the Initialize Tape (INZTAP) command. You use the field to write the owner identifier of the volume or to write information about the contents of the volume. If the identifier is less than 14 bytes, the system justifies the value and pads it with blanks.

- Label Standard Level

Identifies the version of ISO/ANSI standards. For Version 3, the program places a 3 in this field.

ISO/ANSI Standard Data Set Label 1 (HDR1/EOV1/EOF1)

Table 13. Format of the ISO/ANSI Standard Data Set Label 1 (HDR1/EOV1/EOF1)

Offset	Data Type	Length	Contents	File Type
0	Character	3	Label identifier	
3	Character	1	Label number	
4	Character	17	File Identifier	
21	Character	6	File Set Identifier	
27	Character	4	File Section Number	
31	Character	4	File Sequence Number	
35	Character	4	Generation Number (not used)	
39	Character	2	Version Number (not used)	
41	Character	6	Creation Date	
47	Character	6	Expiration Date	
53	Character	1	Accessibility	
54	Character	6	Block Count, Low Order (Trailer labels only)	
60	Character	13	System Code (Trailer labels only)	
73	Character	7	Reserved	

The ISO/ANSI standard data set label 1 (HDR1/EOV1/TRL1) is 80 characters in length and is used to identify each data set. The system records the data set label in ASCII.

The manual describes the contents and function of each of the fields below.

- Label identifier

The characters identify the type of data set label.

HDR Header Label (the beginning of a data set)

EOF Trailer Label (the end of a data set)

EOV Trailer Label (the end of a data set that is continued on another volume)

- Label number

The relative position of this label within the set of labels of the same type; it is always 1 for the data set label 1.

- File Identifier

A unique identification code to identify the data set (file). If the data set ID is less than 17 characters, it is left justified and padded with blanks.

- File Set Identifier

This field contains the volume identifier from the volume labels of the first volume in a multivolume data set.

- File Section Number

This field contains the relative volume number of this volume in a multivolume data set. The field is 0001 for a single volume data set.

- File Sequence Number

  This field indicates the relative position of the data set within a multiple data set group. The value will be in EBCDIC displayable characters for 0001–9999. For numbers larger than 9999 that will not fit in the 4–character field as an EBCDIC displayable character set, the first byte will a '?' ('6F'x for EBCDIC labels. The last three bytes will be a binary number from 1 to 64000.

- Generation Number (not used)

  If the data set is a part of a generation data group, this field contains a number that indicates a absolute generation number. OS/400 does not use this field.

- Version Number (not used)

  If the data set is part of a generation data group, this field contains a number indicating the version number of the generation. OS/400does not use this field.

- Creation Date

  The creation date of the data set. the program shows the date in the format cyyddd, where:

  c = century (blank=19; 0=20; 1=21; and so on)

  yy = year (00-99)

  ddd = day of the current year (001-366)

  Note that the century code gives the first two digits of the year, not the actual century. For example, a blank which translates into 19 indicates a year in the 1900s, not in the nineteenth century.

- Expiration Date

  The program considers the date in which the data set expires. Expired refers to a data set being allowed to be overwritten. The user specifies the date in the open tape file used in writing the tape. The OS/400 ignores the expiration date on input, and verifies the expiration date on output. The date is in the format cyyddd, where:

  c = century (blank=19; 0=20; 1=21; and so on)

  yy = year (00-99)

  ddd = day of the current year (001-366)

  Note that the century code gives the first two digits of the year, not the actual century. For example, a blank which translates into 19 indicates a year in the 1900s, not in the nineteenth century.

- Accessibility

  A code number indicating the security status of the data set. OS/400 does not use this field. The following values indicate data set security values created by other systems:

  blank No data set access protection

  1 Password protection (Required for read/write/deletion)

  3 Password protection (Required for write/deletion)

  other character
  Protected volume, no access possible

- Block Count (Trailer labels only)

  The field in the trailer labels contains the number of data blocks in the data set on the current volume. The field in the header labels contains hex zeros.

- System Code (Trailer labels only)

  A unique code that identifies the system that created the data set.

IBMOS400
> IBM OS/400

IBM OS/VS 370
> IBM MVS operating system

ISO/ANSI Standard Data Set Label 2 (HDR2/EOV2/EOF2)

Table 14. Format of the IBM Standard Data Set Label 2 (HDR2/EOV2/EOF2)

Offset	Data Type	Length	Contents
0	Character	3	Label identifier
3	Character	1	Label number
4	Character	1	Record Format
5	Character	5	Block Length
10	Character	5	Record Length
15	Character	35	Reserved for Operating System
50	Character	2	Buffer offset
52	Character	28	Reserved

The ISO/ANSI standard data set label 2 (HDR2/EOV2/TRL2) is 80 characters in length and is used to identify additional information about the data set. The system records the data set label in ASCII.

The manual describes the contents and function of each of the fields below.

- Label identifier

  The characters identify the type of data set label.

  HDR   Header Label (the beginning of a data set)

  EOF   Trailer Label (the end of a data set)

  EOV   Trailer Label (the end of a data set that is continued on another volume)

- Label number

  The relative position of this label within the set of labels of the same type; it is always 2 for the data set label 2.

- Record Format

  An alphabetic character that indicates the format of the records in the data set. The record format indicates to the operating system how to interpret the blocks of data that the program reads from the tape volume.

  F   Fixed length records

  V   Variable length records

  U   Undefined length records

- Block Length

  A number indicating the block length, in bytes, of the data blocks on the tape. The number in this field can range from 18 to 2048. This block length should include the buffer offset and padding.

  Note that the 18-byte to 2048 byte limit on block length is an ISCII/ASCII standard. You may specify larger blocks (up to 9999 bytes) with the agreement of the interchange parties. However, for tapes with Version 3 labels, exceeding the 2048 byte limit may create incompatible tapes.

  Interpretation of the number depends on the associated record format field as follows:

Record format F
> Maximum block length

Record format D
> Maximum block length that includes the 4 byte length indicator in the records and the optional block prefix.

Record format S
> Maximum block length that includes the optional block prefix, plus one or more pairs of 5 byte segment control words and segments.

- Record Length

  A number indicating the record length, in bytes, of the logical records on the tape volume. The interpretation of the number depends on the Record Format field.

  – Record format F

    Actual record length

  – Record format D

    Maximum record length that includes the 4 byte length indicator in the records

  – Record format S

    Maximum record length that excludes all the 5 byte segment control words that describe the record. If the record length is larger than 99999, this field is 0.

- Reserved for Operating System

  The content of this field is optional for each operating system. OS/400 has chosen similar meaning to the same bytes in the IBM Standard Data Set Label 2.

  – Tape Density (1 byte)

    3   1600 bpi

    4   3200 bpi

    5   6250 bpi

    blank All other densities/formats

  – Data Set Position (1 byte)

    A code indicating a volume switch is as follows:

    0   No volume switch has occurred

    1   A volume switch previously occurred

  – Job/Job Step Identification (17 bytes)

    This field identifies the job and job step that created or extended the data set. OS/400 does not use this field.

  – Tape Recording Technique (2 bytes)

    This field indicates the tape recording technique used in creating the data set.

    blank No Improved Data Recording Capability (IDRC) used.

    'P '   Improved Data Recording Capability (IDRC) used.

  – Control Character (1 byte)

    A printer control codeindicating whether the program uses a control character set to create the data set and the type of control characters the program uses:

    A   Contains ISO/ANSI control letters

    blank Contains no control characters

  – Buffer Alignment Block (1 byte)

    OS/400 does not use this field.

- – Block Attribute

  A code indicating the block attribute used to create the data set:

  B   Blocked records

  blank   Records not blocked

- – Reserved (11 bytes)

  OS/400 does not use this field.

- • Buffer Offset

  The length of the block prefix (from 0 to 99). Used to determine the length of an optional prefix that may be a part of a physical block on tape. The version of the prefix for variable and spanned record formats is known as a block descriptor word (BDW). A BDW is always 4 bytes long and contains the block length of the physical record it describes, including the BDW.

ISO/ANSI Standard User Labels (UHL and UTL)

Table 15. Format of the ISO/ANSI Standard User Labels (UHL and UTL)

Offset	Data Type	Length	Contents
0	Character	3	Label identifier
3	Character	1	Label number
4	Character	76	User data

The ISO/ANSI standard user labels (UHL and UTL) is 80 bytes in length and is used to identify additional information about the data set. All information in the user labels is user data. The user labels directly follow the header/trailer group and can be up to any number of user labels per data set.

The program describes the contents and function of each of the fields below.

- • Label identifier

  The characters identify the type of data set label.

  UHL   User Header Label (the beginning of a data set)

  UTL   User Trailer Label (the end of a data set)

- • Label number

  The relative position of this label within the set of labels of the same type; there is a no limit of user labels per data set.

- • User Data

  This field specifies any user information in the user label. The system ignores the information in this set but passes it to a specified user label program.

Other ISO/ANSI Labels

ISO/ANSI standards also allow up to 9 volume labels, up to 9 header and trailer labels, but OS/400 will not create these ″extra″ labels. OS/400 will read and ignore these labels if created on another operating system.

I/O Feedback Area

The system communicates the results of I/O operations to the program using Operating System/400 (OS/400) messages and I/O feedback information. The system updates the I/O feedback area for every I/O operation unless your program uses a blocked record I/O. In that case, the system updates the feedback area only when you read or write a block of records. Some of the information reflects the last

record in the block. Other information, such as the count of I/O operations, reflects the number of operations on blocks of records and not the number of records. See your high-level language manual to determine if your program uses blocked record I/O.

The I/O feedback area consists of two parts: a common area and a file-dependent area. The file-dependent area varies by the file type.

Common I/O Feedback Area

Table 16. Common I/O Feedback Area

Offset	Data Type	Length	Contents
0	Binary	2	Offset to file-dependent feedback area.
2	Binary	4	Write operation count. Updated only when a write operation completes successfully. For blocked record I/O operations, this count is the number of blocks, not the number of records.
6	Binary	4	Read operation count. Updated only when a read operation completes successfully. For blocked record I/O operations, this count is the number of blocks, not the number of records.
10	Binary	4	Write-read operation count. Updated only when a write-read operation completes successfully.
14	Binary	4	Other operation count. Number of successful operations other than write, read, or write-read. Updated only when the operation completes successfully. This count includes update, delete, force-end-of-data, force-end-of-volume, change-end-of-data, release record lock, and acquire/release device operations.
18	Character	1	Reserved.
19	Character	1	Current operation.

hex 01 Read or read block or read from invited devices

hex 02 Read direct

hex 03 Read by key

hex 05 Write or write block

hex 06 Write-read

hex 07 Update

hex 08 Delete

hex 09 Force-end-of-data

hex 0A Force-end-of-volume

hex 0D Release record lock

hex 0E Change end-of-data

hex 0F Put delete

hex 11 Release device

hex 12 Acquire device

20	Character	10	Not applicable to tape and diskette.
30	Character	2	Device class:

Byte 1:

hex 00	Database
hex 01	Display
hex 02	Printer
hex 04	Diskette

Table 16. Common I/O Feedback Area (continued)

Offset	Data Type	Length	Contents	
			hex 05	Tape
			hex 09	Save
			hex 0B	ICF
			Byte 2 (if byte 1 contains hex 00):	
			hex 00	Nonkeyed file
			hex 01	Keyed file
			Byte 2 (if byte 1 does not contain hex 00 and contains either hex 04 or hex 05):	
			hex 08	Spooled
			hex 1A	9347 Tape Unit
			hex 1B	9348 Tape Unit
			hex 1C	9331-1 Diskette Unit
			hex 1D	9331-2 Diskette Unit
			hex 2A	6346 Tape Unit
			hex 2B	2440 Tape Unit
			hex 2C	9346 Tape Unit
			hex 2D	6331 Diskette Unit
			hex 2E	6332 Diskette Unit
			hex 3A	3430 Tape Unit
			hex 3B	3422 Tape Unit
			hex 3C	3480 Tape Unit
			hex 3D	3490 Tape Unit
			hex 49	6367 Tape Unit
			hex 4A	6347 Tape Unit
			hex 4E	6341 Tape Unit
			hex 4F	6342 Tape Unit
			hex 50	6133 Diskette Unit
			hex 53	6366 Tape Unit
			hex 54	7208 Tape Unit
			hex 5A	6343 Tape Unit
			hex 5B	6348 Tape Unit
			hex 5C	6368 Tape Unit
			hex 64	6344 Tape Unit
			hex 65	6349 Tape Unit
			hex 66	6369 Tape Unit
			hex 67	6380 Tape Unit
			hex 68	6378 Tape Unit
			hex 69	6390 Tape Unit
			hex 70	6379 Tape Unit
			hex 71	9331-11 Diskette Unit
			hex 72	9331-12 Diskette Unit
			hex 73	3570 Tape Unit
			hex 74	3590 Tape Unit
			hex 75	6335 Tape Unit
			hex 76	1/4-inch Cartridge Tape[1]
			hex 77	1/2-inch Cartridge Taped[1]
			hex 78	1/2-inch Reel Tape[1]
			hex 79	8mm Cartridge Tape[1]

Table 16. Common I/O Feedback Area (continued)

Offset	Data Type	Length	Contents
32	Character	10	Device name. The name of the device for which the operation just completed. Supplied only for display, printer, tape, diskette, and ICF files. For printer or diskette files being spooled, the value is *N. For ICF files, the value is the program device name. For other files, the value is the device description name.
42	Binary	4	Length of the record processed by the last I/O operation (supplied only for an ICF, display, tape, or database file). On ICF write operations, this is the record length of the data. On ICF read operations, it is the record length of the record associated with the last read operation.
46	Character	80	Reserved.
126	Binary	2	Number of records retrieved on a read request for blocked records or sent on a write or force-end-of-data or force-end-of-volume request for blocked records. Supplied only for database, diskette, and tape files.
128	Binary	2	Record length. For output, the field value is the record format length, including first-character forms control, option indicators, source sequence numbers, and program-to-system data. If the value is zero, use the field at offset 42.
			For input, the field value is the record format length, including response indicators and source sequence numbers. If the value is zero, use the field at offset 42.
130	Character	2	Reserved.
132	Binary	4	Current block count. The number of blocks of the tape data file already written or read. For tape files only.
136	Character	8	Reserved.

Note:

1. If the device is not listed above, use one of the following general categories:
    - hex 76—1/4″ Cartridge Tape
    - hex 77—1/2″ Cartridge Tape
    - hex 78—1/2″ Reel Tape
    - hex 79—8mm Cartridge Tape

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation

81

Software Interoperability Coordinator
3605 Highway 52 N
Rochester, MN 55901-7829
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming Interface Information

This publication is intended to help you to program for tape and diskette input devices, and tape and diskette output devices. This publication documents Product-Sensitive Programming Interface and Associated Guidance Information provided by Operating System/400.

Product-Sensitive programming interfaces allow the customer installation to perform tasks such as diagnosing, modifying, monitoring, repairing, tailoring, or tuning of this IBM software product. Use of such interfaces creates dependencies on the detailed design or implementation of the IBM software product. Product-Sensitive programming interfaces should be used only for these specialized purposes. Because of their dependencies on detailed design and implementation, it is to be expected that programs written to such interfaces may need to be changed in order to run with new product releases or versions, or as a result of service.

Programming Interface and Associated Guidance Information is identified where it occurs, either by an introductory statement to a chapter or section or by the following: **Product-Sensitive Programming Interface:**

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

400
Application System/400
AS/400
BRMS/400
COBOL
IBM
ILE C
ILE RPG
Operating System/400
OS/400

Other company, product, and service names may be the trademarks or service marks of others.

—

The following AS/400 manuals contain information you may need.The manuals list their title and base order number.

- *Backup and Recovery*, SC41-5304-03, provides the system programmer with information on planning and accomplishing a backup and recovery strategy and how to recover from disk unit failures and disaster. It includes information about setting up and managing:
  - Journaling, access path protection, and commitment control
  - User auxiliary storage pools (ASPs)
  - Disk protection (device parity, mirrored, and checksum)

  It also includes topics on advanced backup and recovery, such as save-while-active support, saving and restoring to a different release, and programming tips and techniques.

- *ICF Programming*, SC41-5442-00, provides the application programmer with information to write application programs that use AS/400 communications and ICF files. It also contains information on data description specifications (DDS) keywords, system-supplied formats, return codes, file transfer support, and program examples.*Communications Configuration*, SC41-5401-00, provides information about configuration commands and defining lines, controllers, and devices.

- *DDS Reference*, SC41-5712-01, provides the application programmer with detailed descriptions of entries and keywords to describe:
  - Database files (both logical and physical).
  - Certain device files (for displays, printers, and ICF) external to the user's programs.

- *Data Management*, SC41-5710-00, provides the application programmer with information about using files in application programs. It describes fundamental structure and concepts of data management support on the AS/400 system, copying files, and temporarily changing files when running an application program through overrides and file redirection.

- *DB2 UDB for AS/400 Database Programming*, SC41-5701-02, discusses, for the application programmer or system programmer, the AS/400 database organization, including information on how to create, describe, and manipulate database files on the system.

- *Local Device Configuration*, SC41-5121-00, provides the system operator or system administrator with information on how to do an initial local hardware configuration and how to change that configuration. It also contains conceptual information for device configuration, and planning information for device configuration on the 9406, 9404, and 9402 System Units.

- *Distributed Data Management*, SC41-5307-00, provides the application programmer or system programmer with information about remote file processing. It describes:
  - How to define a remote file to OS/400 DDM (distributed data management)
  - How to create a DDM file
  - File utilities that are supported through DDM
  - The requirements of OS/400 DDM as related to other systems

- *Application Display Programming*, SC41-5715-00, provides information about creating and maintaining screens for applications, creating online help information, and working with display files on the AS/400 system.

- *Printer Device Programming*, SC41-5713-03, provides information on:
  - How to understand and control printing
  - Printing elements and concepts
  - Printer file support
  - Support for print spooling
  - Printer connectivity
  - Advanced function printing
  - Printing with personal computers

- *CL Programming*, SC41-5721-02, provides a wide-ranging discussion of programming topics, including:
  - A general discussion of objects and libraries
  - Control language (CL) programming
  - Controlling flow and communicating between programs
  - Working with objects in CL programs
  - Creating CL programs
  - Predefined and immediate messages and message handling
  - Defining and creating user-defined commands and menus

85

- – Application testing that is including:
    - Debug mode
    - Breakpoints
    - Traces
    - Display functions
- *CL Reference (Abridged)*, SC41-5722-03, provides a description of the control language (CL) and its commands. Each command definition includes syntax diagram, parameters, default values, and keywords.
- *Automated Tape Library Planning and Management*, SC41-5309-02, provides information about tasks performed with an automated tape library (ATL).The citation describes recommended methods for ATL design and use.
- *System API Programming*, SC41-5800-00, provides conceptual and guidance information for a new user of the OS/400 application programming interface (APIs).
- *System API Reference*, SC41-5801-03, provides information for experienced programmers on how to use the application programming interfaces (APIs) to such OS/400 functions as:
    - – Dynamic Screen Manager
    - – Message handling
    - – National language support
    - – Problem management
    - – Registration facility
    - – Security
    - – Tape management exit programs
    - – Work management
- *Security - Reference*, SC41-5302-03, provides the system programmer (or someone who is assigned the responsibilities of a security officer) with information about:
    - – System security concepts
    - – Planning for security
    - – Setting up security on the system
- *System Operation*, SC41-4203-00, provides information about how:
    - – To use the system unit control panel and console
    - – To send and receive messages
    - – To respond to error messages
    - – To start and stop the system
    - – To use control devices

- – To work with program temporary fixes (PTFs)
- – To process and manage jobs on the system

Index

Numerics

A

B

C

Readers' Comments — We'd Like to Hear from You

AS/400e
Tape and Diskette Device Programming
Version 4

Publication No. SC41-5716-01

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you?   ☐ Yes   ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

_____   _____
Name   Address

_____   _____
Company or Organization

_____
Phone No.

Readers' Comments — We'd Like to Hear from You

SC41-5716-01

IBM ®

Fold and Tape   **Please do not staple**   Fold and Tape
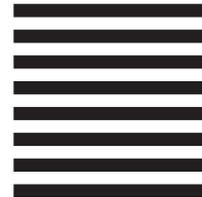
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL   PERMIT NO. 40   ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM CORPORATION
ATTN DEPT 542 IDCLERK
3605 HWY 52 N
ROCHESTER MN 55901-7829

Fold and Tape   **Please do not staple**   Fold and Tape

SC41-5716-01

IBM ®

Spine information:

IBM   AS/400e   OS/400 Tape and Diskette Device Programming V4R4   Version 4