DFSORT/VSE

General Information

Version 3 Release 4

GC26-7039-03

DFSORT/VSE

General Information

Version 3 Release 4

GC26-7039-03

Note!

Before using this information and the product it supports, be sure to read the general information under “Notices” on page v.

Fourth Edition (May 1998)

This edition replaces and makes obsolete the previous edition, GC26-7039-02. The technical changes for this edition are summarized under “Summary of Changes,” and are indicated by a vertical bar to the left of a change. A vertical bar to the left of a figure caption indicates that the figure has changed. Editorial changes that have no technical significance are not noted.

This edition applies to Version 3 Release 4 of DFSORT/VSE, Program Number 5746-SM3, and to any subsequent releases until otherwise indicated in new editions or technical newsletters. Make sure you are using the correct edition for the level of the product.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

Comments may be addressed to:

International Business Machines Corporation
RCF Processing Department
G26/050
5600 Cottle Road
SAN JOSE, CA 95193-0001
U.S.A.

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

Contents

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, programs, or services, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter 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
500 Columbus Avenue
Thornwood, NY 10594
U.S.A.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

DFSORT
ECKD
ESA/370
ESA/390
IBM
VSE/ESA

Preface

If you are considering buying DFSORT/VSE, you should read all of this book to learn how DFSORT/VSE can help you process and arrange information.

DFSORT/VSE General Information briefly explains what DFSORT/VSE is and how you can benefit from using it. This book describes how DFSORT/VSE sorts, merges, and copies small to very large amounts of data. It also describes how DFSORT/VSE can be used to analyze your data, produce reports, and process your two-digit year data.

About This Book

You should read this book if you are a data processing manager, application programmer, system programmer, or analyst evaluating DFSORT/VSE for use in your organization.

Chapter 1, "What is DFSORT/VSE?" on page 1, provides an overview of DFSORT/VSE. It also describes the improvements in the latest release of DFSORT/VSE.

Chapter 2, "What Are the Benefits of DFSORT/VSE?" on page 7, describes the benefits of DFSORT/VSE that can help you to meet your ever-changing business needs. It describes the ways in which DFSORT/VSE can provide increased efficiency, flexibility and productivity in your business.

Chapter 3, "What Can You Do with DFSORT/VSE?" on page 15, describes some of the tasks that you can accomplish using DFSORT/VSE.

DFSORT/VSE Publications

DFSORT/VSE General Information is part of a more extensive DFSORT/VSE library. These books can help you use DFSORT/VSE more effectively. The books in the library are listed below.

Task	Publication	Order Number
Application programming	*DFSORT/VSE Application Programming Guide*	SC26-7040
Planning for, installing, customizing, and tuning DFSORT/VSE	*DFSORT/VSE Installation and Tuning Guide*	SC26-7041
Learning to use DFSORT/VSE	*Getting Started with DFSORT/VSE*	SC26-7101
Diagnosing failures and interpreting messages	*DFSORT/VSE Messages, Codes and Diagnosis Guide*	SC26-7132
Quick reference	*DFSORT/VSE Reference Summary*	SX26-6008

You can order a complete set of DFSORT/VSE publications with the order number SBOF-6130, except for *DFSORT/VSE Licensed Program Specifications* (GC26-7038), which must be ordered separately.

DFSORT/VSE Library Softcopy Information

A softcopy version of the DFSORT/VSE library is available on the CD-ROM shown in the table that follows. The *IBM Online Library VSE Collection* contains all of the DFSORT/VSE books for Releases 2, 3, and 4, with the exception of the *DFSORT/VSE Reference Summary*, and books from other VSE libraries.

Order Number	Title
SK2T-0060	*IBM Online Library VSE Collection*

DFSORT/VSE on the World Wide Web

For news, tips, and examples, visit the DFSORT/VSE home page at URL:

`http://www.ibm.com/storage/dfsortvse/`

Chapter 1. What is DFSORT/VSE?

Throughout the years, the VSE operating system has met the information technology needs of your business. With the advent of ESA technology, VSE has become even more powerful. As the amount of data processing on VSE/ESA continues to grow, you will need powerful, reliable tools to help you process your data quickly and easily.

Sorting is a necessary part of day-to-day processing. Since it is used frequently and consumes a large amount of resources, you need the right sort product to help you manage and organize your data. To meet this need, IBM proudly offers you DFSORT/VSE Version 3 Release 4.

DFSORT/VSE provides fast and efficient sorting, merging, copying, reporting, and analysis of your business information. DFSORT/VSE runs on the VSE/ESA operating system and is as effective for small amounts of data and simple jobs, such as alphabetizing a list of names, as it is for very large amounts of data and complex jobs, such as running a billing system.

Overview of DFSORT/VSE

DFSORT/VSE offers three basic functions — sorting, merging, and copying. These functions provide the basis for a host of features offered in DFSORT/VSE. Using DFSORT/VSE, you can:

- Sort using data space and GETVIS areas.
- Sort using virtual disk space rather than DASD.
- Specify secondary allocation amounts for SAM ESDS work files which allows more efficient use of work space.
- Sort, merge, and transform a wide variety of dates with two-digit years according to a specified sliding or fixed century window.
- Create reports quickly and easily using DFSORT/VSE's ICETOOL utility. ICETOOL also enables you to do a variety of utility-type functions such as data analysis and validation.
- Specify subsets of data using the SKIPREC and STOPAFT options.
- Control which records to keep in the final output file by using the INCLUDE and OMIT control statements.
- Rearrange or eliminate fields in your record by using the INREC and OUTREC control statements.
- Sum numeric information into totals or eliminate duplicate records by using the SUM control statement.
- Control the use of DFSORT/VSE's STXIT exit routine for abend recovery processing.
- Invoke DFSORT/VSE from an application program written in a high level language such as COBOL, PL/I, or RPG II with the Auto-Report Feature.
- Use EBCDIC or ISCII/ASCII collating sequences, specify your own collating sequence, or use locale processing.

- Call user-written exit routines to expand the capabilities of DFSORT/VSE.

What's New in DFSORT/VSE Version 3 Release 4?

DFSORT/VSE Version 3 Release 4 continues IBM's strategy of providing performance and productivity improvements. These improvements help your applications run more efficiently and provide additional features that can simplify your work. They are described in more detail in the subsections that follow.

Performance

DFSORT/VSE provides these performance improvements:

- Enhanced data processing methods and input/output processing
- Extended Count Key Data (ECKD) support
- User control of VSAM buffers
- Work file enhancements

Enhanced Data Processing Methods and Input/Output Processing

DFSORT/VSE provides improved data processing methods and input/output processing for:

- Incore and non-incore dataspace sorting
- Incore and non-incore getvis sorting
- Non-incore partition sorting
- Copy and merge applications

Extended Count Key Data (ECKD) Support

DFSORT/VSE provides improved ECKD disk device support for input, output, and work files by using the ECKD command set.

User Control of VSAM Buffers

The VSAMBSP installation option allows users to control the number of buffers DFSORT/VSE can use for VSAM (or SAM ESDS accessed as VSAM) input and output file processing.

Work File Enhancements

DFSORT/VSE provides the following work file processing enhancements:

- All work files are now closed at the end of an application.
- Additional work file extents can now be used, if available, when end of extent is encountered regardless of whether STXIT is in effect.
- All extents of an SD work file can now be used instead of only the first extent.

Productivity

DFSORT/VSE provides these productivity improvements:

- Additional Year 2000 formats
- Additional OUTREC features
- INCLUDE/OMIT enhancements
- New ZDPRINT feature
- New Online Message Explanations (OME) feature

Additional Year 2000 Formats

You can now use the new Y2S and Y2B formats for SORT, MERGE and OUTREC, in addition to the previously available Y2C, Y2Z, Y2P, and Y2D formats. The new formats, like the old formats, can be used in conjunction with the Y2PAST installation and run-time option to handle two-digit year data.

Y2S Format: The new Y2S format works like the previously available Y2C/Y2Z formats, but does special handling of binary zeros, blanks, and binary ones.

For SORT and MERGE, Y2S interprets *real* two-digit character or zoned decimal year data according to the century window, while collating binary zeros, blanks, and binary ones in their *normal* positions relative to the *real* years.

For OUTREC, Y2S transforms *real* two-digit character or zoned decimal year data to *real* four-digit character year data according to the century window, while transforming binary zeros, blanks, and binary ones to an equivalent four-byte form.

For example, you can use these control statements:

```
OPTION Y2PAST=1915
SORT FIELDS=(1,2,Y2S,A)
OUTREC FIELDS=(1,2,Y2S)
```

to sort and transform two-digit year character input data of:

```
98, binary ones, 14, 00, 15, blanks, binary zeros
```

to four-digit year character output data of:

```
binary zeros, blanks, 1915, 1998, 2000, 2014, binary ones
```

Y2B Format: The new Y2B format handles two-digit binary year data.

For SORT and MERGE, Y2B interprets two-digit year binary year data according to the century window.

For OUTREC, Y2B transforms two-digit binary year data to four-digit character year data according to the century window.

For example, you can use these control statements:

```
OPTION Y2PAST=1915
SORT FIELDS=(1,1,Y2B,A)
OUTREC FIELDS=(1,1,Y2B)
```

to sort and transform two-digit binary year input data of:

```
X'62', X'00', X'0E', X'1A', X'0F'
```

to four-digit year character output data of:

```
1915, 1926, 1998, 2000, 2014
```

Additional OUTREC Features

OUTREC has new operands that allow you to use the new features described in the subsections that follow.

Lookup and Change (p,m,lookup): OUTREC now has a very useful feature called lookup and change. It tells DFSORT/VSE to look at the value of a field in your input record and match it to the values you have set up in a table. When a match is found, the associated string in the table is substituted for the original value in the output record. This makes it easy for you to substitute meaningful words or phrases for cryptic values (for example, *FERN* can be changed to *FERNALL BROS.*).

Hexadecimal Display (p,m,HEX and p,HEX): OUTREC now allows you to convert data to printable hexadecimal characters (EBCDIC).

Editing Masks (p,m,f,edit): OUTREC now allows you to convert numeric fields to printable decimal characters with sign control (for example, *+* or *blank* for positive numbers and - or *()* for negative numbers) and editing control (for example, commas every n digits, decimal point, leading *$*, suppressed or non-suppressed leading zeros, and so on). Twenty-six pre-defined editing masks are available for commonly used numeric editing patterns, encompassing many of the numeric notations used throughout the world. In addition, a virtually unlimited number of numeric editing patterns are available via user-defined editing masks.

Column Support (c:): OUTREC allows you to specify the output column for any field or separator, making it easier to design your output record.

INCLUDE/OMIT Enhancements

You can now:

- Use a significantly larger number of INCLUDE and OMIT conditions
- Include or omit all records using the new ALL or NONE operands

New ZDPRINT Feature

You can now convert positive ZD sums to printable numbers using the new ZDPRINT feature.

Online Message Explanations (OME)

With Online Message Explanations (OME), you request that an explanation of a DFSORT/VSE message be displayed on the console of your display station. OME makes it easier for you to access information about messages you receive and can help you diagnose and correct errors more quickly. You can use the DFSORT/VSE OME just as you would use the VSE/ESA OME.

Additional Enhancements

DFSORT/VSE also provides these additional enhancements:

- STXIT improvements
- New DIAGINF feature
- New NRECOUT feature

STXIT Improvements

The STXIT=MIN installation option and the MINSTXIT run-time option provide a new DFSORT/VSE capability that removes a STXIT-related performance degradation. You can use the STXIT=MIN or MINSTXIT option if you have COBOL or PL/I programs that invoke DFSORT/VSE and use E15/E35 user exit routines to process records.

The IBM-supplied default is changed from STXIT=YES to STXIT=MIN.

New DIAGINF Feature

The new DIAGINF installation option allows you to request diagnostic information (diagnostic messages and a dump), regardless of the options in effect at run time.

This option is useful when DFSORT/VSE is called by an application, such as COBOL, and the COBOL source code is missing or unavailable. In these cases, you are unable to change anything in the application, including the DFSORT/VSE control statements. With the new DIAGINF option, you don't need to change the application. You can force DFSORT/VSE diagnostic messages and a dump if an abend occurs.

New NRECOUT Feature

The new NRECOUT installation and run-time option gives you more control over the way DFSORT/VSE handles the situation when it does not write any records to the output file. You can tell DFSORT/VSE to:

- Continue processing and pass back a return code of 0
- Continue processing and pass back a return code of 4
- Terminate processing and pass back a return code of 16

With this new option, you can control the return code and the outcome of the application when DFSORT/VSE does not write any records to the output file.

Chapter 2. What Are the Benefits of DFSORT/VSE?

This chapter describes the benefits of DFSORT/VSE. In an ever-changing business environment, DFSORT/VSE can help your business by increasing:

- Efficiency
- Flexibility
- Productivity

Efficiency

Performance improvements with DFSORT/VSE help reduce batch window times and increase system throughput. Reduced input/output times and work space requirements are also possible using DFSORT/VSE. It is designed specifically to optimize efficiency and processing speed by:

- Using data space
- Using virtual storage (GETVIS areas)
- Using extended addressing
- Supporting the COBOL/VSE FASTSRT compile-time option

In addition, techniques of DFSORT/VSE can increase efficiency in the day-to-day operation of your business by:

- Using dataspace sorting
- Using getvis sorting
- Expanding system capacity
- Improving system availability
- Exploiting system resources
- Enhancing system space management

Dataspace Sorting

Dataspace sorting is a DFSORT/VSE capability that uses data space available with VSE/ESA systems on ESA/370 and ESA/390 processors in place of intermediate work space for sort applications. A data space is an area of contiguous storage backed by processor or auxiliary storage, whichever is necessary as determined by the system.

With dataspace sorting, data space is used as an extension of partition virtual storage for storing records during a sort. Dataspace sorting reduces CPU time and elapsed time by using data space in place of DASD work files. You can specify dataspace sorting at installation time for global use or at run time for specific jobs.

Dataspace sorting can be used with both fixed-length records (FLR) and variable-length records (VLR). If you currently use partition program area sorting for your applications, using dataspace sorting should provide significant CPU and elapsed time performance improvements.

Getvis Sorting

Getvis sorting is a DFSORT/VSE capability that uses partition virtual storage for sort applications. Getvis sorting is similar to dataspace sorting except that it uses the partition GETVIS area for virtual storage instead of data space. If you currently use partition program area sorting for your applications, using getvis sorting should provide significant CPU and elapsed time performance improvements. Also, using the partition GETVIS area exploits the available partition virtual storage above 16 MB virtual.

Expanding System Capacity

With 31-bit addressing, DFSORT/VSE takes advantage of the additional virtual storage available above 16 MB virtual. 31-bit addressing gives the system more capacity for calling programs and user exits resulting in better performance for DFSORT/VSE and for the overall system. The additional virtual storage means more sort applications can run with less work space, thereby improving elapsed time, CPU time, and start I/Os (SIOs).

Improving System Availability

By moving its program modules above 16 MB virtual, DFSORT/VSE eases the virtual storage constraints on your 24-bit shared virtual area (SVA) and partition program area. With more available storage, your system can handle more sort applications and process data quicker, resulting in better overall system performance and throughput.

Exploiting System Resources

DFSORT/VSE is typically used many times each day at large sites. Because sorting can occupy much of a site's CPU time and I/O and channel resources, even small improvements in performance can translate directly into significant dollar and machine resource savings. DFSORT/VSE processing methods are designed to give significant overall performance benefits while balancing:

- Elapsed time used by the program (elapsed time is always important and is often the primary consideration when evaluating the efficiency of very large sorts)
- CPU time used by the program
- I/O utilization (device connect time, channel usage, and SIOs)

DFSORT/VSE's dataspace sorting capability uses data spaces available on VSE/ESA. By using dataspace sorting, more sort applications can be processed without work space, and, those applications that still require work space can be processed with a smaller amount. In both cases, dataspace sorting can improve your batch window. By avoiding or reducing I/O to DASD, DFSORT/VSE's dataspace sorting improves elapsed time, CPU time, SIOs, and channel usage. When DFSORT/VSE determines that dataspace sorting should be used, the performance benefits are achieved automatically and require no changes to your sort applications.

You can get similar benefits even when there is not enough data space available. If dataspace sorting cannot be used, you can tell DFSORT/VSE to use the getvis sorting feature instead. Getvis sorting can also improve elapsed and CPU time,

decrease SIOs, and reduce work space requirements on systems with ESA/370 and ESA/390 architecture.

Getvis and dataspace sorting provide similar performance, and both allow you to exploit the storage areas above 16 MB virtual for sorting. Getvis sorting uses the storage resources of the partition and is preferable when an application should not depend on the availability of system virtual storage. Dataspace sorting uses global system resources and is preferable when an application can take extensive advantage of system virtual storage.

Since processor storage plays an important role in the use of getvis and dataspace sorting, DFSORT/VSE bases its use of its very efficient sorting methods on the amount of processor storage it can use without causing excessive paging on the system. DFSORT/VSE dynamically controls the storage areas used by dataspace sorting and getvis sorting. DFSORT/VSE only monitors the paging activity of the system when getvis and data space size are set to the maximum value.

Note: When you install DFSORT/VSE, dataspace sorting and getvis sorting are **not** enabled. Use the DSPSIZE installation or run-time option to enable dataspace sorting. Use the GVSIZE installation or run-time option to enable getvis sorting.

Enhancing System Space Management

Even though 31-bit addressing, dataspace sorting, and getvis sorting reduce work space requirements, some sort applications with extremely large files may still require the use of work space. DFSORT/VSE provides for the secondary allocation of SAM ESDS work files with the VSE/VSAM Space Management for SAM Feature. This means that sort processing will continue even for cases where the primary allocation has been exhausted. By allowing DFSORT/VSE to calculate and allocate the secondary work space, you will no longer need to do precise primary allocation calculations yourself, thus saving you valuable time. This means less work space will be wasted and more jobs will be able to use the allocated space. Secondary work allocation also means fewer jobs will terminate due to an incorrect primary work space allocation and, therefore, fewer jobs will need to be re-run. You can see how better space management can help you increase your system throughput.

If you have a File Management System installed at your site and set the DFSORT/VSE FMS installation option to YES, the File Management System provides additional assistance with space management. File Management Systems provide benefits such as:

- Dynamic logical and physical device assignment
- Dynamic primary and secondary extent allocation
- Output file truncation of unused space

Flexibility

DFSORT/VSE offers you greater flexibility by providing you with facilities for:

- National Language Support
- Customizing DFSORT/VSE
- Invoking DFSORT/VSE

- Invoking the ICETOOL utility

National Language Support

With DFSORT/VSE, you can define the cultural environment for sorting, merging, including, and omitting records. The cultural environment includes such items as the native language of the user, the character set for a language, and the collating rules for a language. You use the LOCALE option at installation or run time to select the cultural environment. Then, DFSORT/VSE can collate and compare your information according to the rules of the locale you select.

Using this National Language Support feature, you can:

- Obtain results based on the cultural and local characteristics of a language you select
- Obtain results for your cultural environment without costly pre- and post-processing
- Process data for another cultural environment by selecting a different locale

In addition, DFSORT/VSE facilitates the translation of its messages into different languages; flexible positioning of variables within a message is supported. And, DFSORT/VSE is shipped with both English and Japanese DFSORT/VSE messages and Online Message Explanations (OME).

Customizing DFSORT/VSE

You can customize the way you use DFSORT/VSE, and how it handles your data, with:

- Installation and run-time options
- User exit routines
- Alternate collating sequences

Installation and Run-Time Options

DFSORT/VSE installation defaults can be customized to suit your needs. The ILUINST macro can be used to modify the IBM-supplied DFSORT/VSE installation defaults. Many of these installation options have an impact on performance, including:

- Virtual storage limit
- Data space limit
- GETVIS area limit

At run time, you can temporarily override most specifications made during installation by specifying:

- Control statements in SYSIPT for JCL-invoked applications
- Control statements in the parameter list for program-invoked applications

User Exit Routines

For selected jobs, you can direct DFSORT/VSE to transfer control to your own routines at predesignated points (user exits). User exit routines extend the flexibility of DFSORT/VSE beyond the many standard capabilities it provides, allowing you to:

- Open files
- Initialize files
- Insert, delete, or alter records in files
- Terminate DFSORT/VSE
- Summarize records
- Close files
- Complete housekeeping tasks

Like the other routines written at your site, user exit routines must follow standard linkage conventions. Routines can be written in any language that can pass the location or address of a record or parameter list in register 1. Assembler and PL/I are examples. PL/I users, however, must use the special PL/I SORT subroutine facilities.

In addition, user-written routines and programs that invoke DFSORT/VSE can:

- Reside above or below 16 MB virtual
- Run in 24-bit or 31-bit mode
- Pass data that resides above or below 16 MB virtual to DFSORT/VSE

Alternate Collating Sequences

DFSORT/VSE usually determines the order of the output records using the standard EBCDIC or ISCII/ASCII collating sequence. However, you can use locale processing, the ALTSEQ installation option, or the ALTSEQ control statement to alter the EBCDIC collating sequence.

Invoking DFSORT/VSE

With DFSORT/VSE, you have the flexibility of choosing how to invoke DFSORT/VSE processing. Depending on the needs of your business, you can invoke DFSORT/VSE processing in the following ways:

- With an EXEC job control statement in the input stream using SORT as the program name. This means that DFSORT/VSE is not initiated from another program but is *directly invoked*.
- With a program written in either assembler language, COBOL, PL/I, or RPG II with the Auto-Report Feature. This means that DFSORT/VSE is initiated from another program and is *program invoked*.

Invoking the ICETOOL Utility

As with DFSORT/VSE, you have the flexibility of choosing how to invoke the ICETOOL utility. Depending on the needs of your business, you can invoke ICETOOL either directly or from a program. ICETOOL allows operator statements (and comments) to be supplied in a SYSIPT file or in a parameter list passed by a calling program. For each operator supplied in the parameter list, ICETOOL puts

information in the parameter list pertaining to that operation, thus allowing the calling program to use the information derived by ICETOOL.

Productivity

With DFSORT/VSE, your productivity is improved because you can:

- Prepare for the turn of the century by correctly processing your two-digit year data
- Take advantage of File Management Systems
- Create reports and analyze data using the ICETOOL utility

Preparing for the Year 2000

DFSORT/VSE's Year 2000 features will help you prepare for the turn of the century by correctly processing your two-digit year data. The Y2C, Y2Z, Y2P, Y2D, Y2S, and Y2B formats, in conjunction with the Y2PAST installation and run-time option, allow you to handle two-digit year data in the following ways:

- Set the appropriate century window for your applications (for example, 1917-2016 or 1950-2049).
- Order two-digit character, zoned decimal, packed decimal, decimal, or binary year data according to the century window using SORT or MERGE (for example, order 96 representing 1996 before 00 representing 2000 in ascending sequence, or order 00 before 96 in descending sequence).
- Transform two-digit character, zoned decimal, packed decimal, decimal, or binary year data to four-digit character (or zoned decimal) year data according to the century window using OUTREC (for example, transform 96 to 1996 and 00 to 2000).

The PD0 format allows you to order parts of packed decimal fields, such as month and day in date fields, using SORT or MERGE.

The PZ, PSI, and ZSI formats allow you to transform packed decimal and zoned decimal fields, such as month and day in date fields, to character fields using OUTREC. The character and hexadecimal string separators allow you to insert literals in reformatted fields using OUTREC (for example, '/' in mm/dd/yyyy fields).

Interaction with File Management Systems

You can specify that DFSORT/VSE should attempt to take advantage of the benefits provided by the File Management System installed at your site. These benefits could include:

- Dynamic logical and physical device assignment
- Dynamic primary and secondary extent allocation
- Output file truncation of unused space

Just set the FMS installation option to YES or NO, depending on whether or not you have a File Management System installed at your site.

Reporting and Analyzing Data

With ICETOOL, a versatile DFSORT/VSE utility, you can do reporting and analysis of data at your site. ICETOOL allows you to perform multiple operations on one or more files in a single job step. This batch front-end utility uses the capabilities of DFSORT/VSE to perform the operations you request.

ICETOOL has thirteen operators: COPY, COUNT, DEFAULTS, DEFINE, DISPLAY, MODE, OCCUR, RANGE, SELECT, SORT, STATS, UNIQUE, and VERIFY. By using one operator or a combination of these operators, you can easily create applications that perform a variety of tasks including:

- Sorting one or more input files to one or more output files
- Creating multiple copies of input files
- Creating output files containing subsets of input files based on various criteria
- Creating detailed reports allowing control of title, date, time, page numbers, headings, lines per page, field formats, and total, maximum, minimum, and average values
- Creating reports showing unique values for selected character and numeric fields and the number of times each occurs
- Creating reports or output files for records with: duplicate values, non-duplicate values, or values that occur n times, less than n times, or more than n times
- Creating a wide variety of reports using the preferred date, time, and numeric notations of individual countries
- Creating a report showing the DFSORT/VSE installation defaults in use
- Printing messages that give statistical information for selected numeric fields such as minimum, maximum, average, total, count of values, and count of unique values
- Printing messages that identify invalid decimal values
- Using three different modes, (stop, continue, and scan) to control error checking and actions after error detection for groups of operators

ICETOOL can be called directly or from a program. It also produces messages and return codes describing the results of each operation and any errors detected. Although you generally do not need to look at the DFSORT/VSE messages produced as a result of an ICETOOL run, they are available if you need them.

As an example of what you can do with ICETOOL, Figure 1 on page 14 shows possible ICETOOL statements and DLBL statements you can use for an ICETOOL run. The DEFINE operators specify the format and length of the input records. The SORT operator selects books only from publishers VALD and WETH and sorts them by publisher and title. The DISPLAY operator prints a report with title and headers from the sorted subset output created by the SORT operator.

```
//DLBL BKIN,'SORT.SAMPIN',,VSAM,CAT=VSESPUC
//DLBL BKADD,'SORT.SAMPADD',,VSAM,CAT=VSESPUC
//DLBL DAPUBS,'%%SORT',,VSAM,RECORDS=100,RECSIZE=173,CAT=VSESPUC
  ...
* Specify the format and record length of the input file
 DEFINE NAME(BKIN) TYPE(F) LENGTH(173)
* Select books from VALD and WETH and sort them
 SORT FROM(BKIN,BKADD) TO(DAPUBS) USE
  USTART
    SORT FIELDS=(106,4,A,1,30,A),FORMAT=CH
    INCLUDE COND=(106,4,EQ,C'VALD',OR,106,4,EQ,C'WETH'),FORMAT=CH
  UEND
* Specify the format and record length of the input file
 DEFINE NAME(DAPUBS) TYPE(F) LENGTH(173)
* Print a report showing the sorted books
 DISPLAY FROM(DAPUBS) LIST(LST) BLANK -
   TITLE('Books from VALD and WETH') PAGE DATE(DM4/)
   HEADER('Book Title') ON(1,30,CH) -
   HEADER('Publisher') ON(106,4,CH)
```

Figure 1. Example of an ICETOOL Run

Figure 2 shows an example of a report that can be produced by these operators. You can see how easy it is to create a report that contains a sorted subset of the records in input files using ICETOOL.

```
Books from VALD and WETH  - 1 -  29/05/1998

Book Title                      Publisher
------------------------------  ---------
CELLS AND HOW THEY WORK         VALD
COMPLETE SPANISH DICTIONARY     VALD
EDITING SOFTWARE MANUALS        VALD
FREUD'S THEORIES                VALD
INTRODUCTION TO BIOLOGY         VALD
NOVEL IDEAS                     VALD
SHORT STORIES AND TALL TALES    VALD
STRATEGIC MARKETING             VALD
VIDEO GAME DESIGN               VALD
ZEN BUSINESS                    VALD
ANTICIPATING THE MARKET         WETH
CIVILIZATION SINCE ROME FELL    WETH
COMPUTERS: AN INTRODUCTION      WETH
EIGHTEENTH CENTURY EUROPE       WETH
GUIDE TO COLLEGE LIFE           WETH
GUNTHER'S GERMAN DICTIONARY     WETH
REBIRTH FROM ITALY              WETH
SYSTEM PROGRAMMING              WETH
THE INDUSTRIAL REVOLUTION       WETH
```

Figure 2. Example of an ICETOOL Report

Chapter 3. What Can You Do with DFSORT/VSE?

DFSORT/VSE is a highly useful data processing tool. With DFSORT/VSE, you can:

- Prepare for the Year 2000 and beyond
- Rearrange records in a file with the sort function
- Combine files with the merge function
- Duplicate files with the copy function
- Produce reports using the ICETOOL utility
- Analyze your data using the ICETOOL utility
- Process files with control statements

Preparing for the Year 2000 and Beyond

If you use applications that make use of years represented by only two digits, you are probably anticipating problems with these applications in the year 2000. With its Year 2000 features, DFSORT/VSE now provides you with two-digit year data ordering and transformation capabilities as follows:

- Data formats (Y2C, Y2Z, Y2P, Y2D, Y2S, and Y2B) allow the correct ordering of two-digit year data with SORT and MERGE along with transformation of two-digit year data to four-digit year data with OUTREC.
- An installation and run-time option, Y2PAST, allows you to set a sliding or fixed century window that defines the correct collation order for Y2 format data.
- A data format, PD0, allows ordering of parts of PD fields, such as month and day, with SORT and MERGE.
- Data formats, PZ, PSI, and ZSI, allow transformation of PD and ZD fields, such as month and day, with OUTREC.
- Character and hexadecimal strings can be used to specify separators for date fields with OUTREC.

For example, if you set the sliding century window to 81 and the current year is 1998, the century window will be 1917 to 2016. DFSORT/VSE will interpret Y2 format data for SORT, MERGE, and OUTREC as follows:

Figure 3. Sliding Century Window: 1917-2016

Two-digit Year Data	Interpretation
00	2000
16	2016
17	1917
61	1961
62	1962
99	1999

Or, if you set the fixed century window to 1962, the century window will be 1962 to 2061. DFSORT/VSE will interpret Y2 format data for SORT, MERGE, and OUTREC as follows:

Figure 4. Fixed Century Window: 1962-2061

Two-digit Year Data	Interpretation
00	2000
16	2016
17	2017
61	2061
62	1962
99	1999

Notice the differences in how the same data is interpreted, based on how the century window is set.

Rearranging Records in a File

The primary function of DFSORT/VSE is sorting files, rearranging records from one sequence to another. The SORT control statement activates the sort function. You can use additional control statements to:

- Include or exclude records. You can create subsets of files based on relational conditions that:
  - Compare two fields
  - Compare a field and a constant
  - Search for a constant within a field value
  - Search for a field value within a constant
  - Test the state (on or off) of selected bits in a field
  - Compare fields and constants based on EBCDIC sequence, ISCII/ASCII sequence, a sequence you define, or locale processing

  You can combine relational conditions with AND and OR operators to create complex logical expressions.
- Reformat records before or after other processing. You can edit, delete, and rearrange fields, align fields in columns, and insert strings, blanks, or zeros before, between, or after fields.
- Sum numeric values in your records. You can use DFSORT/VSE's summing capabilities to produce totals or subtotals, or summaries of records based on combinations of numeric or character control fields. You can also use these capabilities to eliminate duplicate records.
- Change the sequence for control fields. The sequence for each control field can be ascending or descending. DFSORT/VSE can use the standard EBCDIC sequence, the ISCII/ASCII sequence, a sequence that you define, or locale processing.

- Skip records. You can use the SKIPREC option to specify the number of records you want to skip before starting to sort the input files.
- Accept records. You can use the STOPAFT option to specify the maximum number of records you want accepted for sorting.
- Preserve the original sequence of records from input to output for sort applications when control fields are the same.

Combining Files

With DFSORT/VSE's merge function, you can combine up to nine previously sorted input files. The MERGE control statement activates the merge function. As with a sort, you can also use additional control statements to:

- Include or exclude records
- Reformat records after other processing
- Sum numeric values in your records or eliminate duplicate records
- Change the sequence for control fields
- Preserve the original sequence of records from input to output for merge applications when the control fields are the same

Duplicating Files

With DFSORT/VSE's copy function, you can:

- Create an output file that duplicates the information in your input files
- Specify copying using either the SORT or MERGE control statement
- Use additional control statements to include or exclude records and reformat records after other processing
- Specify the number of records you want to skip before starting to copy the input files using the SKIPREC option
- Specify the maximum number of records you want accepted before starting to copy the input files using the STOPAFT option

Producing Reports with ICETOOL

With DFSORT/VSE's ICETOOL utility, you can produce reports quickly and easily. You can produce detailed reports with section headings and statistics, page headings and statistics, and a summary page with headings and statistics. You can produce reports that identify and count unique values, duplicate values, or values that occur a specific number of times.

ICETOOL can even create a report showing the DFSORT/VSE installation defaults selected at your site.

Analyzing Data with ICETOOL

With DFSORT/VSE, you can use the ICETOOL utility to perform data analysis. You can create multiple copies or create different subsets of the input files in a single job step. You choose the selection criteria that gives you the information you need. For example, you can decide to:

- Keep the first or last record with each unique field
- Keep the records with duplicate values
- Keep all the records with values that occur a specific number of times
- Keep all the records that have values within a certain range

With ICETOOL, you can also identify and display invalid decimal values in a readable format and show the location of the record in the file. You can conditionally process a group of dependent operators; if an operation fails, the remaining operators are not processed. Or, you can use ICETOOL to unconditionally process a group of independent operators; each operator is processed even though other operators succeed or fail.

Processing Files with Control Statements

With DFSORT/VSE control statements, you direct how your files are processed. You can use control statements to provide information about the control fields and collating sequences, and to describe the input data. In addition, you can specify whether the order of equally collating records is to be preserved from input to output, and whether any user-written routines are to be included at run time.

The following are required control statements in a DFSORT/VSE application:

SORT Describes control fields and the number of input and work files for a sort application, or specifies a copy application. Indicates whether you want ascending or descending order for the sort.

MERGE Describes control fields and the number of input files for a merge application, or specifies a copy application. Indicates whether you want ascending or descending order for the merge.

RECORD Describes record length and type information.

Other control statements give you more control over how your files are processed:

ALTSEQ Modifies the default collating sequence that was set at installation time.

ANALYZE Specifies that DFSORT/VSE is to analyze the control statements, perform optimization calculations, issue appropriate messages, and stop the run without actually sorting, merging, or copying.

INCLUDE Specifies that only those records whose fields meet certain criteria are included.

INPFIL Describes the input files and specifies the procedures that are to be followed when an input file is opened and closed.

INREC Specifies how records are reformatted before they are processed.

MODS Specifies the use of one or more user exit routines in a DFSORT/VSE application.

OMIT Specifies that any records whose fields meet certain criteria are omitted.

OPTION Overrides installation defaults and supplies optional information.

OUTFIL Describes the output file and specifies the procedures that are to be followed when an output file is opened and closed.

OUTREC Specifies how records are reformatted before they are output.

SUM Specifies that numeric summary fields in records with equal control fields are to be summarized in one record and that the other records are to be deleted. Also specifies that duplicate records are to be eliminated.

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

DFSORT/VSE
General Information
Version 3 Release 4

Publication No. GC26-7039-03

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
GC26-7039-03

Cut or Fold
Along Line

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

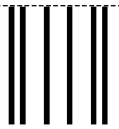

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

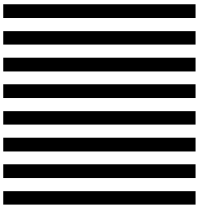

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
RCF Processing Department
G26/050
5600 Cottle Road
SAN JOSE, CA 95193-0001

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

Cut or Fold
Along Line

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

DFSORT/VSE
General Information
Version 3 Release 4

Publication No. GC26-7039-03

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
GC26-7039-03

Cut or Fold
Along Line

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

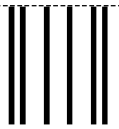

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

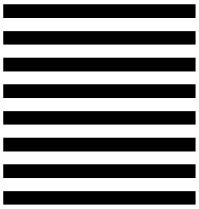

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
RCF Processing Department
G26/050
5600 Cottle Road
SAN JOSE, CA 95193-0001

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

Cut or Fold
Along Line

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC26-7039-03