

DFSORT: Summary of Changes by Release

DFSORT Web Site

For papers, online books, news, tips, examples and more, visit the DFSORT website at URL:

<http://www.ibm.com/storage/dfsor>

Contents

DFSORT: Summary of Changes by Release	1
Introduction	1
z/OS 3.2 - Base (September, 2025)	1
z/OS 3.1 - Base (September, 2023)	1
z/OS V2R5 - Base (September, 2021)	1
IBM Integrated Accelerator for z Sort	1
z/OS V2R3 and z/OS V2R4 - PTFs (September, 2020)	1
IBM Integrated Accelerator for z Sort	1
z/OS V2R4 - Base (September, 2019)	1
Regular Expressions	1
Unicode comparisons	1
ASCII free format numeric	1
Encryption support	2
Message changes	2
New Messages	2
New Reserved Words for Symbols	2
z/OS V2R2 and z/OS V2R3 - PTFs (November, 2018)	2
Exploit zHPF for SORTWORK datasets	2
z/OS V2R3 - Base (September, 2017)	2
Sort/Merge of Unicode Transformation Format (UTF8/UTF16/UTF32) data	2
z/OS V2R1 and z/OS V2R2 - PTFs (September, 2016)	2
Performance Improvements	3
z/OS V2R1 and z/OS V2R2 - PTFs (March, 2016)	3
Transfer of blocks of records between DFSORT and user E15/E35 exits.	3
z/OS V2R2 - Base (September, 2015)	3
High Performance Ficon (HPF) Exploitation	3
Date functions	3
Message changes	3
z/OS V2R1 - Base (September, 2013)	4
64-bit support	4
Performance and resource usage improvements	4
XTIOT, uncaptured UCBs and DSAB above 16 megabytes virtual	4
Extended address volumes	4
Alphanumeric comparison tests	5
More PARSE fields	5
Repeating PARSE fields	5
Alphanumeric tests for PARSE fields	5
Insert a string at the end of variable-length records	6
Symbol enhancements	6
Improved diagnostics	6
z/OS V1R10 and z/OS V1R12 - PTFs (October, 2010)	6
RESIZE	6
Update the trailer record with OUTFIL	6
Date Field Arithmetic	7
SET and PROC symbols in control statements	7
Accept n records with OUTFIL	7
New translation functions	7
Begin group when key changes	8
Larger header/trailer fields for OUTFIL	8
RC8 and RC12 for COUNT	8
More fields for DISPLAY	8
Larger fields for ICETOOL	8

Longer reports for DISPLAY/OCCUR	8
Timestamp with microseconds	8
z/OS V1R12 - Base (September, 2010)	8
Improvements in Performance	9
Improved Diagnostics	9
File size for Improperly Closed VSAM Data Sets	9
Improved First Failure Data Capture	9
Dynamic Allocation Improvements	9
Extended Address Volumes	9
XTIOT, Uncaptured UCBs and DSAB Above 16 Megabytes Virtual	10
z/OS V1R5 and z/OS V1R10 - PTFs (November, 2009)	10
JOINKEYS Application	10
Date Field Conversions	10
MERGE Operator	11
MERGEIN Alternate ddnames	11
New 24-Bit Parameter List Codes	11
z/OS V1R10 - Base (September, 2008)	11
Installation Options	11
VSAM Extended Address Volume	11
Improvements in Performance and Resource Usage and Reporting	12
Syntax for Run-Time Parameters	12
Improved Diagnostics	12
z/OS V1R5 - PTFs (July, 2008)	12
Find and replace	12
Group operations	13
DATASORT	13
SUBSET	13
SELECT first n duplicates	14
SPLICE with non-blank fields	14
DISPLAY with count	14
DISPLAY/OCCUR with multiple and multipart titles	14
DISPLAY/OCCUR without carriage control	14
COUNT in output record	14
COUNT with add/subtract	15
BLKSIZE default for input DUMMY	15
SKIP=0L default for SECTIONS	15
SORTOUT=ddname default for FNAMES	15
z/OS V1R5 and Release 14 - PTFs (April, 2006)	15
INCLUDE and OMIT Enhancements	15
INREC and OUTREC Enhancements	16
OUTFIL Enhancements	17
Symbol Enhancements	18
ICETOOL Enhancements	19
SORT and MERGE Enhancements	19
SUM Enhancements	19
Other Enhancements	19
z/OS V1R5 and Release 14 - PTFs (December, 2004)	19
ICETOOL Enhancements	19
OUTFIL Enhancements	20
INREC and OUTREC Enhancements	21
SORT and MERGE Enhancements	22
INCLUDE and OMIT Enhancements	23
SUM Enhancement	23
Other Enhancements	23

z/OS V1R5 - Base (March, 2004)	23
Improvements in Performance	23
Installation and Run-Time Options	23
Release 14 - PTFs (February, 2003)	24
ICETOOL Enhancements	24
OUTFIL Enhancements	24
INREC and OUTREC Enhancements	25
INCLUDE and OMIT Enhancements	25
SORT, MERGE and SUM Enhancement	25
Other Enhancements	25
Release 14 - PTFs (March, 2002)	26
z900	26
Multiple Hiperspaces	26
Managed Tapes	26
ICETOOL Enhancements	26
OUTFIL Enhancements	27
INREC and OUTREC Enhancements	28
SUM Enhancement	28
SORT and MERGE Enhancements	29
INCLUDE and OMIT Enhancements	29
Other Enhancements	29
Release 14 - PTFs (July, 2000)	29
New Device Support	29
Larger Tape Block Sizes with OS/390 R10	30
DFSMSrmm-Managed Tapes	30
VSAM Processing	30
HFS Files	30
Productivity and Easier Migration from other Sort Products	31
New Generation of Year 2000 Features	32
Release 14 - Base (September, 1998)	32
Symbols for Fields and Constants	32
Improvements in Performance, Capacity and Storage Usage	32
Time-of-Day Option Controls	33
Repackaging	33
OUTFIL Processing Enhancements	33
ICETOOL Enhancement	33
Installation and Run-Time Option Enhancements	33
Other Enhancements	34
OS/390 and MVS/ESA Only	34
Release 13 - PTFs	34
New Device Support	34
Year 2000 Features	34
OS/390 Registration	35
Performance Improvements	35
Floating Point for SUM	36
Security Improvements	36
EXCPVR Processing Removed	36
Release 13 - Base (May, 1995)	36
DFSORT's Performance Booster for The SAS System	36
Dynamic Hipersorting	36
Performance	36
OUTFIL Processing	36
National Language Support	37
ICETOOL Enhancements	37

INCLUDE/OMIT Substring Search	38
SMF Type-16 Record Enhancements	38
Other Enhancements	38
Release 12 - PTFs	39
New Device Support	40
Other Enhancements	40
Release 12 - Base (September, 1992)	40
Performance	40
MVS/ESA and MVS/XA Only	40
ICETOOL	40
Floating Sign Format	41
Panels	41
Other Enhancements	41
Release 11.1 - PTFs	42
New Device Support	42
Other Enhancements	42
Release 11.1 - Base (February, 1991)	43
ICETOOL	43
Panels	44
Other Enhancements	44
Performance	44
Release 11 - PTFs	45
New Device Support	45
Other Enhancements	45
Release 11 - Base (November, 1988)	45
Enhancements	45
Release 10 - Base (January, 1988)	46
Enhancements	46
Release 9 - Base (April, 1987)	47
Enhancements	47
MVS/ESA and MVS/XA only	48
Release 8 - Base (March, 1986)	48
Enhancements	48
Release 7.1 - Base (June, 1985)	48
Enhancements	48
Release 7 - Base (January, 1985)	48
New Device Support	48
Other Enhancements	48
Release 6 - Base (January, 1984)	49
Enhancements	49

DFSORT: Summary of Changes by Release

Introduction

For reference and historical perspective, this document shows the changes to each DFSORT Release by PTFs, and at the base level. These lists were adapted from the *DFSORT Application Programming Guide* for the various Releases.

Changes made by PTFs to a lower Release (for example, V2R5) are always incorporated into the higher Release (for example, 3.1) either in the base or with PTFs.

z/OS 3.2 - Base (September, 2025)

No changes in content.

z/OS 3.1 - Base (September, 2023)

No changes in content.

z/OS V2R5 - Base (September, 2021)

IBM Integrated Accelerator for z Sort

Enable DFSORT to leverage new on chip sort accelerator known as the IBM Integrated Accelerator for Z Sort to provide CPU and Elapsed time improvements without the need for application changes.

z/OS V2R3 and z/OS V2R4 - PTFs (September, 2020)

IBM Integrated Accelerator for z Sort

Enable DFSORT to leverage new on chip sort accelerator known as the IBM Integrated Accelerator for Z Sort to provide CPU and Elapsed time improvements without the need for application changes.

z/OS V2R4 - Base (September, 2019)

Regular Expressions

DFSORT now supports Regular expressions in the following DFSORT comparison operands: COND, INCLUDE, OMIT, BEGIN, END, WHEN and TRILID. This support allows DFSORT users to use Regular expressions in their batch jobs for additional filtering capabilities. Regular expressions contain a series of characters that define a pattern of text to be matched, which allows for more robust filtering capabilities.

Unicode comparisons

DFSORT now supports the usage of Unicode data formats (UTF-8, UTF-16 and UTF-32) in the following comparison operands: COND, INCLUDE and OMIT.

ASCII free format numeric

DFSORT now supports two new ASCII free format numeric data formats (AUF and ASF) in SORT and MERGE operands. The new ASCII free format numeric data formats (AUF and ASF) are also now supported in the following comparison operands: COND, INCLUDE and OMIT.

Encryption support

DFSORT now supports reading and writing to sequential extended format data sets, VSAM extended format data sets (KSDS, ESDS, RRDS, VRRDS) and PDSE data sets.

Message changes

The following messages are new, changed, or no longer issued for z/OS DFSORT Installation and Customization in z/OS V2R4.

The text for existing DFSORT messages has been changed. Any automated actions based on the presence of these messages should be evaluated.

• ICE000I

• ICE201I

New Messages

The following message is new.

• ICE167A

New Reserved Words for Symbols

The following are new DFSORT/ICETOOL reserved words which are no longer allowed as symbols: AUF and ASF. If you used any of these words as a symbol previously, you must change them. For example, if you used ASF, you can change it to asf.

z/OS V2R2 and z/OS V2R3 - PTFs (November, 2018)

Exploit zHPF for SORTWORK datasets

Enable DFSORT to create channel programs that exploit System z High Performance FICON (zHPF) with SORTWORK datasets to provide I/O performance improvements without the need for application changes.

z/OS V2R3 - Base (September, 2017)

Sort/Merge of Unicode Transformation Format (UTF8/UTF16/UTF32) data

DFSORT now supports the Sorting and Merging of Unicode data. This support allows users to:

- SORT/MERGE Unicode Data with control field length of 1 to 450 Unicode characters for UTF-8 format data.
- SORT/MERGE Unicode Data with control field length of 1 to 450 Unicode characters for UTF-16 format data.
- SORT/MERGE Unicode Data with control field length of 1 to 450 Unicode characters for UTF-32 format data.

z/OS V2R1 and z/OS V2R2 - PTFs (September, 2016)

Performance Improvements

Improved performance for several DFSORT functions that generate code at run time. DFSORT's new performance improvements delivered by PTFs UI40658 and UI40659 will improve performance in both CPU and elapsed times.

z/OS V2R1 and z/OS V2R2 - PTFs (March, 2016)

Transfer of blocks of records between DFSORT and user E15/E35 exits.

DFSORT exit processing of E15/E35 exits are enhanced to support the transfer of blocks of records between DFSORT and E15/E35 exits. With this new support for inserting blocks of records it will expand the functionality of the existing E15/E35 exits which will significantly reduce the number of calls of the E15/E35 exits, and will reduce the transfer of records between user's storage and DFSORT storage. Further-more, the new E15/E35 block support does not affect the existing support of single record interface.

z/OS V2R2 - Base (September, 2015)

High Performance Ficon (HPF) Exploitation

A program that invokes DFSORT, ICETOOL or ICEGENER will now be able to exploit the High Performance Ficon (zHPF) hardware. When zHPF is available, DFSORT SORT or MERGE applications will use BSAM for SORTIN, SORTOUT and OUTFIL data sets.

Date functions

Two new date functions to calculate the week number of a given date and Age as date duration are added.

WEEKNUM function converts a given Julian/Gregorian date to number of week. There are 2 versions of this function, the standard USA format and the ISO format.

- i WEEKNUM=USA function returns an integer in the range of 1 to 54 that represents the week of the year. The week starts with Sunday, and January 1 is always in the first week.
- i WEEKNUM=ISO function returns an integer in the range of 1 to 53 that represents the week of the year. The week starts with Monday and includes 7 days. Week 1 is the first week of the year to contain a Thursday, which is equivalent to first week containing January 4.

AGE function returns a date duration that specifies the number of years, months, and days between an input date and current date.

- i Age=YMD produces an 8-byte result which has duration in years (0-9999), months(00-12), and days (00-31).
- i Age=YM produces a 6-byte result which has duration in years (0-9999), months(00-12).
- i Age=YD produces a 7-byte result which has duration in years (0-9999), days(00-366).

Message changes

The text for existing DFSORT messages ICE000I and ICE288I has been changed. Any automated actions based on the presence of these messages should be evaluated.

z/OS V2R1 - Base (September, 2013)

64-bit support

Eligible user programs and exits can now be written to:

- Call DFSORT from a program in 64-bit addressing mode (AMODE 64) using a new 64-bit parameter list and the entry name ICEMAN64 or SORT64
- Use DFSORT E15, E35 and E32 exits running in 64-bit addressing mode (AMODE 64)
- Pass 64-bit addressed records to DFSORT using new 64-bit parameter lists for E15, E32 and E35 exits.

Performance and resource usage improvements

DFSORT has been enhanced to improve use of central storage in relation to system activity. The existing EXPMAX, EXPOLD and EXPRES installation options will now be evaluated in conjunction with available resources at run-time which can give you better control over how DFSORT uses available central storage resources. The likelihood of over committed central storage resources and excessive paging has been reduced, which can provide improved reliability and performance for all workloads including DFSORT.

The shipped default for EXPOLD is changed from MAX to 50%. The shipped to system default for EXPRES is changed from 0 to 10%.

DFSORT has been enhanced to allocate storage in smaller increments and then check resource availability before each additional increment is reserved. This can allow concurrent sort applications to better share available storage and allows sorts to stop allocating additional storage if resources are no longer available. Additionally, DFSORT's dynamic allocation of work data sets will be adjusted to reduce the likelihood of failure when a sort application is unable to obtain all of the expected central storage and must then use more disk work space than expected.

A new TUNE installation default allows you to specify whether DFSORT should allocate storage in increments with additional disk work space to minimize the risk of failure, or to allocate all storage at initialization so disk work space allocation can be reduced.

Blockset merge applications can now use storage above 16 MB virtual with more functions (such as E61, INREC, OUTREC, INCLUDE, OMIT and SUM), providing improved performance and virtual storage constraint relief.

XTIOT, uncaptured UCBs and DSAB above 16 megabytes virtual

Work data sets dynamically allocated by DFSORT will use options for XTIOT, uncaptured UCBs, and DSAB above 16 megabyte virtual to the extent that z/OS supports them.

Extended address volumes

DFSORT support for EAS-eligible data set types on Extended Address Volumes has been enhanced to increase the maximum size of a DFSORT work data set. With full track blocking, the maximum number of tracks that can be used for a single work data set has been increased from 1048576 to 16777215. In situations where DFSORT must use a reduced block size for the work data sets, less than 16777215 tracks can be used.

Alphanumeric comparison tests

New UC, LC, MC, UN, LN and MN keywords (similar to the previously available NUM keyword) now allow you to test a field for various combinations of alphanumeric characters or non-alphanumeric characters using binary (BI) format.

The new keywords allow you to select any of the following alphanumeric character sets:

- ï UC: Uppercase characters (A-Z)
- ï LC: Lowercase characters (a-z)
- ï MC: Mixed case characters (A-Z, a-z)
- ï UN: Uppercase and numeric characters (A-Z, 0-9)
- ï LN: Lowercase and numeric characters (a-z, 0-9)
- ï MN: Mixed case and numeric characters (A-Z, a-z, 0-9)

You can use these new keywords in the following comparison operands: COND, INCLUDE, OMIT, BEGIN, END, WHEN and TRLID.

More PARSE fields

You can now use up to 1000 parsed fields (%0-%999) with the PARSE function; the previous limit was 100 parsed fields (%0-%99).

Repeating PARSE fields

REPEAT=v is a new PARSE option that can be used to repeat a particular parse field definition multiple times.

REPEAT=v can be used with % to specify v identically defined consecutive parsed fields to be ignored.

REPEAT=v can be used with %n, %nn or %nnn to specify v identically defined consecutive parsed fields for which data is to be extracted. The parsed fields will start with the %n, %nn or %nnn field you select and be incremented by one for each repeated parsed field.

Alphanumeric tests for PARSE fields

STARTAFT=an, STARTAT=an, ENDBEFR=an and ENDAT=an can now be used with the PARSE function to start or end when a character from any of various alphanumeric character sets is found. New keywords for an allow you to select any of the following alphanumeric character sets:

- ï UC: Uppercase characters (A-Z)
- ï LC: Lowercase characters (a-z)
- ï MC: Mixed case characters (A-Z, a-z)
- ï UN: Uppercase and numeric characters (A-Z, 0-9)
- ï LN: Lowercase and numeric characters (a-z, 0-9)
- ï MN: Mixed case and numeric characters (A-Z, a-z, 0-9)
- ï NUM: Numeric characters (0-9)

Insert a string at the end of variable-length records

VLTRAIL=string is a new OUTFIL option that allows you to insert a character string (C'string') or hexadecimal string (X'yy...yy') at the end of each variable-length OUTFIL output record.

Symbol enhancements

DFSORT Symbols can now be used for many more operands. The following operands of the form KEYWORD=n can now be specified as KEYWORD=symbol where symbol represents an equivalent number (for example, if you have New_Length,25 in SYMNAMES, you can use LENGTH=New_Length wherever you can use LENGTH=25): ABSPOS, ACCEPT, ADDPOS, AVGLEN, DO, ENDPOS, ENDREC, FIXLEN, ID, IFOUTLEN, INCR, LENGTH, LINES, MAXLEN, RECORDS, REPEAT, SAMPLE, SEQ, SKIPREC, SPLIT1R, SPLITBY, START, STARTPOS, STARTREC, STOPAFT and SUBPOS.

A symbol for a number can now be used for the length of the output field with the CHANGE operand, and for the length of the sequence number with the SEQNUM operand.

A symbol can be used for any of the new %000-%999 parsed fields.

A symbol can be used for string with the new VLTRAIL=string operand.

Improved diagnostics

DFSORT now provides specific reason codes and associated documentation to aid in diagnosing and correcting errors associated with messages ICE083A and ICE141A.

The text for message ICE118I has been changed to recommend the use of FILSZ=En.

The ICE236I options-in-effect message now includes values for TUNE, EXPMAX, EXPOLD and EXPRES.

The text for message ICE285A has been changed for clarification.

z/OS V1R10 and z/OS V1R12 - PTFs (October, 2010)

RESIZE

RESIZE is a new operator of ICETOOL that allows you to create a larger record from multiple shorter records, or multiple shorter records from a larger record. RESIZE makes it easy to "resize" fixed length records; you just indicate the desired length of the output records and DFSORT automatically resizes the records based on the input length.

Update the trailer record with OUTFIL

IFTRAIL is a new OUTFIL option that allows you to update count and total values in an existing trailer (last) record based on the current data records.

Various options of IFTRAIL allow you to identify the trailer record (TRLID), indicate the count and total values to be updated in the trailer record (TRLUPD), and indicate if the first record is a header record (HD=YES).

IFTRAIL makes it easy to update count and total values in an existing trailer record when you add, delete or modify data records.

DFSORT symbols can be used for fields and constants in TRLID. DFSORT symbols can be used for columns and fields in TRLUPD.

Date Field Arithmetic

New options for BUILD and OVERLAY allow you to perform various types of arithmetic on date fields in either Julian or Gregorian form, with 2-digit or 4-digit years, in CH, ZD and PD format.

ADDDAYS, ADDMONS and ADDYEARS can be used to add days, months or years to a date field.

SUBDAYS, SUBMONS and SUBYEARS can be used to subtract days, months or years from a date field.

DATEDIFF can be used to calculate the number of days between two date fields.

NEXTDday can be used to calculate the next specified day of the week for a date field.

PREVDday can be used to calculate the previous specified day of the week for a date field.

LASTDAYW, LASTDAYM, LASTDAYQ and LASTDAYY can be used to calculate the last day of the week, month, quarter or year for a date field.

DFSORT symbols can be used for fields and constants in date arithmetic.

SET and PROC symbols in control statements

SET and PROC symbols can now be included in DFSORT symbols, which in turn can be used in DFSORT and ICETOOL control statements. Up to ten special symbols (JP0-JP9), which can incorporate SET and PROC symbols, as well as other text and system symbols, can be included in the EXEC PARM string when DFSORT is called directly from JCL (PGM=SORT or PGM=ICEMAN) or ICETOOL is called directly from JCL (PGM=ICETOOL). These JPn symbols will be treated as if they were specified in the SYMNAMES data set. This makes it easy to use DFSORT symbols incorporating SET and PROC symbols in DFSORT and ICETOOL control statements.

Accept n records with OUTFIL

ACCEPT=n is a new OUTFIL option that allows you to limit the number of input records accepted for OUTFIL processing. A record is "accepted" if it is not deleted by STARTREC, ENDREC, SAMPLE, INCLUDE or OMIT processing. ACCEPT makes it easy to limit the number of OUTFIL input records selected for processing.

New translation functions

New TRAN=keyword options in BUILD and OVERLAY provide additional character translation capabilities as follows: ASCII characters to EBCDIC characters (TRAN=ATOE); EBCDIC characters to ASCII characters (TRAN=ETOA); binary values to EBCDIC hexadecimal values (TRAN=HEX); EBCDIC hexadecimal values to binary values (TRAN=UNHEX); binary values to EBCDIC bit values (TRAN=BIT) and EBCDIC bit values to binary values (TRAN=UNBIT). This makes it easy to do such translations on specific fields or entire records.

DFSORT symbols can be used for fields in translation functions.

Begin group when key changes

KEYBEGIN=(p,m) is a new option for WHEN=GROUP clauses that allows a new group to begin whenever a specified key value changes. KEYBEGIN operates in a similar way to the existing BEGIN option, but allows the change in a specified key field to be used as the trigger instead of a logical expression. KEYBEGIN makes it easy to treat records with the same key as a group and propagate fields, identifiers and sequence numbers to the records in each group.

DFSORT symbols can be used for fields in KEYBEGIN.

Larger header/trailer fields for OUTFIL

OUTFIL's HEADERx and TRAILERx options now allow you to specify unedited fields (p,m) up to 32752 bytes (the previous limit was 256 bytes). This makes it easy to use one large header or trailer field rather than multiple header or trailer fields.

RC8 and RC12 for COUNT

ICETOOL's COUNT operator now allows you to use a new RC8 option to set RC=8 (instead of RC=12) or RC=0 based on the count of records in a data set. RC12 can also be used to set RC=12 explicitly instead of by default. RC4, RC8 and RC12 make it easy to set the return code you need for a COUNT operation.

More fields for DISPLAY

ICETOOL's DISPLAY operator now allows you to specify up to 50 HEADER and ON fields (the previous limit was 20). This makes it easy to produce reports that display more information.

Larger fields for ICETOOL

ICETOOL's DISPLAY, OCCUR, SELECT, SPLICE and UNIQUE operators now allow you to specify character ON fields of up to 4000 bytes (the previous limit was 1500 bytes). ICETOOL's DISPLAY and OCCUR operators now allow you to specify hexadecimal ON fields of up to 2000 bytes (the previous limit was 1000 bytes). This makes it easy to use one ON operand for large fields rather than multiple ON operands.

Longer reports for DISPLAY/OCCUR

ICETOOL's DISPLAY and OCCUR operators now allow you to produce reports with a line length up to 8192 bytes (the previous limit was 2048 bytes). This makes it easy to produce reports that display more information.

Timestamp with microseconds

A new DATE5 option can be used to generate a timestamp constant of the form 'yyyy-mm-dd-hh.mm.ss.nnnnnn' with the year, month, day, hours, minutes, seconds and microseconds for the run. This is similar to the previously available DATE4 constant, but includes microseconds (nnnnnn).

z/OS V1R12 - Base (September, 2010)

Improvements in Performance

DFSORT can now use memory object storage as intermediate work space or as an extension of main storage. A memory object is a data area in virtual storage that is allocated above the bar and backed by central storage. Using memory object storage as intermediate work space is the preferred choice.

A new MOWRK installation and run-time option allows you to specify whether memory object storage can be used for intermediate work space or as an extension of main storage, as appropriate, or only as an extension of main storage.

The MOSIZE installation default will still limit the total amount of memory object storage that can be used by a sort application.

Improved Diagnostics

DFSORT now provides additional information in new messages ICE248I and ICE249I when message ICE083A, ICE254I or ICE258I is issued. These new messages will assist with determining why DFSORT was unable to dynamically allocate all of the requested disk work space and what corrective action may be required.

File size for Improperly Closed VSAM Data Sets

When a VSAM data set is opened for output by a program or utility, but not closed properly, the statistics associated with the data set are not updated. This may result in inaccurate file size information for the VSAM data set. When possible, DFSORT will attempt to calculate a more accurate file size for such an improperly closed VSAM fixed-length record data set.

Improved First Failure Data Capture

When an error message (ICExxxA) is issued and the message data set is available, DFSORT will display the diagnostic messages (ICE75xI, ICE8xxI and ICE9xxI) unless MSGPRT=NONE is in effect. This can improve first failure data capture and remove the need to rerun the application to display diagnostic messages.

Dynamic Allocation Improvements

DFSORT's dynamic allocation of work data sets has been enhanced to improve reliability in situations where the disk work space requirements are larger than expected. A new DYNAPCT installation and run-time option allows you to specify additional work data sets to be allocated with zero space. DFSORT only extends these data sets when necessary to complete a sort application. New message ICE278I is issued if these additional work data sets are used.

Extended Address Volumes

DFSORT supports EAS-eligible data set types on Extended Address Volumes to the extent that z/OS supports these data sets. With full track blocking, the maximum number of 1048576 tracks can be used for a single work data set allocated in the cylinder managed space on an Extended Address Volume. Without full track blocking, less than 1048576 tracks may be used.

XTIOT, Uncaptured UCBs and DSAB Above 16 Megabytes Virtual

A program that invokes DFSORT, ICETOOL or ICEGENER can dynamically allocate input, output and work data sets using the options for XTIOT, uncaptured UCBs, and DSAB above 16 megabyte virtual. These data sets are supported to the extent that z/OS supports them.

Note: : A program that invokes DFSORT, ICETOOL or ICEGENER should not allocate other data sets such as message, control, list, count, or symbol data sets using the options for XTIOT, uncaptured UCBs, and DSAB above 16 megabyte virtual. These data set options are not supported and the data set will not be recognized.

z/OS V1R5 and z/OS V1R10 - PTFs (November, 2009)

JOINKEYS Application

A new JOINKEYS application helps you to perform various "join" operations on two files by one or more keys. A JOINKEYS application makes it easy to create joined records in a variety of ways including inner join, full outer join, left outer join, right outer join, and unpaired combinations. The two input files can be of different types (fixed, variable, VSAM, and so on) and have keys in different locations.

Three new control statements can be used for a JOINKEYS application. One JOINKEY statement is required for each input file to indicate the ddname of the file, describe the keys, indicate whether the file is already sorted by those keys, and so on. An inner join is performed by default, but a JOIN statement can be used to specify a different type of join. A REFORMAT statement is used to describe the fields from the two files to be included in the joined records, and optionally an indicator of where the key was found ('B' for both files, '1' for file1 only or '2' for file2 only).

The records from the two input files can be processed in a variety of ways before and after they are joined.

A new JKFROM operand makes it easy to use a JOINKEYS application with a COPY or SORT operator of ICETOOL.

DFSORT symbols can be used for fields and constants specified in the JOINKEYS and REFORMAT statements and in the other DFSORT statements used with a JOINKEYS application.

Date Field Conversions

Y2x, Y4x, TOJUL, TOGREG, WEEKDAY, DT and DTNS keywords for the BUILD and OVERLAY operands make it easy to convert input date fields of one type to corresponding output date fields of another type, and to convert a date field to a corresponding day of the week in several forms.

The date field can be in either Julian or Gregorian form, with 2-digit or 4-digit years, in CH, ZD and PD format. For example, C'ccyyddd', P'yymmdd', Z'dddccyy', C'ccyymmdd', P'dddy', and so on.

CH output date fields can also be displayed with separators. For example, C'ccyy/ddd', C'mm-dd-ccyy', and so on.

The day of the week can be displayed as a 1 digit, 3 character or 9 character constant. For example, C'4', C'WED' or C'WEDNESDAY' for Wednesday.

DFSORT symbols can be used for fields in date conversions.

MERGE Operator

MERGE is a new operator of ICETOOL that allows you to merge up to 50 input files that are already in order by the same keys. MERGE makes it easy to include DFSORT MERGE operations in your ICETOOL jobs.

Various options of MERGE allow you to define the ddnames for the input data sets (FROM), the ddnames for the output data sets (TO), the MERGE and other DFSORT control statements to be used for the MERGE operation (USING), and so on.

DFSORT symbols can be used in the DFSORT statements used with a MERGE operator.

MERGEIN Alternate ddnames

A new MERGEIN operand can be specified on an OPTION statement in DFSPARM or an extended parameter list to supply alternate ddnames for input data sets used for a MERGE application. MERGEIN makes it easy to use any ddnames you like instead of SORTINnn ddnames. You can specify up to 100 ddnames to be merged with MERGEIN.

New 24-Bit Parameter List Codes

DFSORT now accepts and processes the following as 24-Bit parameter list control statement entry codes: X'0E', X'0F' and X'11'.

z/OS V1R10 - Base (September, 2008)

Installation Options

ICEPRMxx members in concatenated PARMLIB can now be used to specify changes to DFSORT's installation options. Each ICEPRMxx member can contain options to be changed for any or all of DFSORT's eight installation environments (JCL, INV, TSO, TSOINV and TD1-TD4). Up to ten ICEPRMxx members can be activated by a START ICEOPT started task command. The options in the activated members will be merged with the ICEMAC defaults at run-time.

A different ICEPRMxx member, or combination of ICEPRMxx members, for different LPARs can be activated at IPL time by including a START ICEOPT command in an appropriate COMMNDxx member in PARMLIB, or at any time by issuing a START ICEOPT command from the console.

ICEPRMxx members are now the recommended way to change DFSORT installation defaults since they are easier to use and more flexible than the old method using the ICEMAC macro and usermods. However, the old method is still supported.

The ICETOOL DEFAULTS operator can be used at any time to produce a report showing the merged PARMLIB/ICEMAC installation default values for each environment that will be used at run-time, as well as the active ICEPRMxx and ICEMAC values.

VSAM Extended Address Volume

DFSORT now supports VSAM input and output data sets on Extended Address Volumes to the extent that z/OS supports these data sets.

Improvements in Performance and Resource Usage and Reporting

Memory object sorting can now be used for DFSORT applications that use the INREC, OUTREC or SUM functions.

The channel programs associated with DFSORT's input, output and work data sets can now reside above 16 megabytes virtual.

Total storage usage for dataspace sorting can now be controlled with the EXPMAX, EXPRES and EXPOLD installation options.

DFSORT now provides new fields ICEMNVLX, ICEMNVLV and ICEMNVLZ in its SMF type 16 record to aid in tuning main storage usage.

The information DFSORT passes to the termination exit (ICETEXIT) now includes setting ICEPTST flag bit 6 on if the phase TCB timings are not valid.

Syntax for Run-Time Parameters

DFSORT now allows additional syntactical forms for all of its PARM and control statement run-time parameters. Any parameter of the form keyword=value can now also be specified as keyword(value) or keyword=(value). Any parameter of the form keyword=(list) can now also be specified as keyword(list).

Improved Diagnostics

DFSORT now provides additional information in messages ICE098I, ICE253I and ICE254I to aid in diagnosing and correcting out-of-space conditions associated with message ICE046A.

DFSORT now provides specific reason codes and associated documentation to aid in diagnosing and correcting errors associated with the following messages: ICE017A, ICE018A, ICE024A, ICE042A, ICE043A, ICE109A, ICE126A and ICE251A.

z/OS V1R5 - PTFs (July, 2008)

Find and replace

FINDREP is a new option that allows you to do various types of find and replace operations on your records. FINDREP makes it easy to replace character or hexadecimal input constants anywhere in your records with character, hexadecimal or null output constants. For input and output constants of different lengths, bytes after the replaced constants will be shifted left or right, as appropriate. For fixed-length records, blanks will be filled in on the right as needed. For variable-length records, the record length will be changed as needed.

FINDREP can be used in an INREC, OUTREC or OUTFIL statement, or in an IFTHEN clause, in the same way BUILD and OVERLAY can be used.

Various options of FINDREP allow you to define one or more input constants and a corresponding output constant (IN, OUT), define one or more pairs of input and output constants (INOUT), start and end the find scan at specified positions (STARTPOS, ENDPOS), stop after a specified number of constants are replaced (DO), increase or decrease the length of the output record (MAXLEN), define the action to be taken if nonblank characters overrun the end of the record (OVERRUN), and specify whether output constants are to replace or overlay input constants (SHIFT).

DFSORT symbols can be used for constants specified with FINDREP.

Group operations

WHEN=GROUP is a new type of IFTHEN clause that allows you to do various types of operations involving groups of records. WHEN=GROUP makes it easy to propagate fields from the first record of a group to the other records of the group, add an identifier to each record of the group, or add a sequence number to each record of the group. These functions are useful by themselves, and can also facilitate other types of group operations such as sorting groups, including or omitting groups, and so on.

WHEN=GROUP can be used in IFTHEN clauses in an INREC, OUTREC or OUTFIL statement in the same way WHEN=INIT can be used.

Various options of WHEN=GROUP allow you to use logical expressions to define the beginning and end of a group (BEGIN, END), define the number of records in a group (RECORDS), and define the fields, identifiers and sequence numbers to be added to the records of each group (PUSH).

DFSORT symbols can be used for columns, fields and constants specified with WHEN=GROUP clauses.

DATASORT

DATASORT is a new operator of ICETOOL that allows you to sort data records between header (first) records and trailer (last) records. DATASORT makes it easy to sort the data records while keeping one or more header records and/or one or more trailer records in place. DATASORT does not require an "identifier" in the header or trailer records; it can treat the first n records as header records and the last n records as trailer records.

Various options of DATASORT allow you to define the number of header records and/or trailer records (HEADER or FIRST, TRAILER or LAST), the ddname for the input data set (FROM), the ddname for the output data set (TO), and the SORT and other DFSORT control statements to be used for the DATASORT operation (USING).

DFSORT symbols can be used for the number of header and trailer records specified with DATASORT.

SUBSET

SUBSET is a new operator of ICETOOL that allows you to create a subset of the input or output records by specifying that you want to keep or remove header (first) records, trailer (last) records, or records with specific relative record numbers. SUBSET makes it easy to keep or remove records based on these criteria. SUBSET does not require an "identifier" or "sequence number" in the records to be kept or removed.

Various options of SUBSET allow you to define the criteria (HEADER or FIRST, TRAILER or LAST, RRN), the ddname for the input data set (FROM), the ddname for the output data set to contain the records that meet the criteria and/or don't meet the criteria (TO, DISCARD), whether the records that meet the criteria are to be kept or removed (KEEP, REMOVE), whether the criteria are to be applied to the input or output records (INPUT, OUTPUT), and DFSORT control statements to be used for the SUBSET operation (USING).

DFSORT symbols can be used for the number of header and trailer records and for the relative record numbers specified with SUBSET.

SELECT first n duplicates

ICETOOL's SELECT operator now allows you to select the first n records with each key or the first n duplicate records with each key. New FIRST(n) and FIRSTDUP(n) options make it easy to select records representing "top" and "bottom" categories (for example, the top 5 students in each class).

DFSORT symbols can be used for n with FIRST(n) or FIRSTDUP(n).

SPLICE with non-blank fields

ICETOOL's SPLICE operator now allows you to create a single record for each key by splicing the base record with every specified nonblank field from each overlay record. A new WITHANY option makes it easy to collect information from multiple records with the same key. You can now do a splice involving duplicate records with nonconsecutive or missing WITH fields, something that could not be accomplished previously with the existing WITHEACH option.

DISPLAY with count

ICETOOL's DISPLAY operator now allows you to display counts in reports. New COUNT('string'), EDCOUNT(formatting), BCOUNT('string') and EDBCOUNT(formatting) options make it easy to print overall record count and break record count statistics in various forms in a report, similar to the existing statistics for a report (overall total, maximum, minimum and average and break total, maximum, minimum and average).

DFSORT symbols can be used for 'string' with COUNT('string') and BCOUNT('string').

DISPLAY/OCCUR with multiple and multipart titles

ICETOOL's DISPLAY and OCCUR operators now allow you to display up to three title lines, each composed of up to three strings. The enhanced TITLE('string1','string2','string3') option makes it easy to use multiple strings for each title, including a combination of inline constants, and constants from DFSORT symbols including system information. The use of up to three TITLE options makes it easy to display multiline titles.

A new TLEFT option allows you to left justify the title lines instead of centering them. A new TFIRST option allows you to only display the title lines on the first page of the report instead of on every page of the report.

DFSORT symbols can be used for 'string1', 'string2' and 'string3' with TITLE('string1','string2','string3').

DISPLAY/OCCUR without carriage control

ICETOOL's DISPLAY and OCCUR operators now allow you to create reports without carriage control characters and with RECFM=FB instead of RECFM=FBA. A new NOCC option makes it easy to suppress the carriage control character. With NOCC, a blank line is used instead of a page eject control character to separate elements of the report.

COUNT in output record

ICETOOL's COUNT operator now allows you to create a count data set with an output record containing the record count. New WRITE(countdd), TEXT('string'), DIGITS(n) and EDCOUNT(formatting) options make it easy to create an output data set with a record containing text and the record count in various forms.

DFSORT symbols can be used for 'string' with TEXT('string').

COUNT with add/subtract

ICETOOL's COUNT operator now allows you to add a value to, or subtract a value from, the record count. New ADD(n) and SUB(n) options make it easy to increase or decrease, respectively, the actual record count to get a resulting modified record count. This is especially useful for dealing with data sets that contain header and/or trailer records.

The resulting modified record count is displayed in the count message in TOOLMSG and in the count data set, and used to determine if the criteria specified by the existing EMPTY, NOTEMPTY, HIGHER(x), LOWER(y), EQUAL(v) or NOTEQUAL(w) option is satisfied.

DFSORT symbols can be used for n with ADD(n) and SUB(n).

BLKSIZE default for input DUMMY

DFSORT will no longer terminate for a SORTIN DD DUMMY or SORTINnn DD DUMMY statement with RECFM and LRECL, but no BLKSIZE. Instead, DFSORT will use an appropriate BLKSIZE to process the DUMMY data set successfully.

Note: If DFSORT's Blockset technique is not selected, DFSORT may still terminate for a SORTIN DD DUMMY or SORTINnn DD DUMMY statement with RECFM and LRECL, but no BLKSIZE.

SKIP=0L default for SECTIONS

DFSORT will no longer terminate when an OUTFIL SECTIONS field is not followed by a keyword (SKIP, HEADER3, TRAILER3). Instead, DFSORT will use a default keyword of SKIP=0L to process the sections successfully with no blank lines between sections on the same page.

DFSORT symbols can be used for section fields.

SORTOUT=ddname default for FNAMES

DFSORT will now use the ddname specified by a SORTOUT=ddname option in DFSPARM, the ddname specified by a SORTOUT=ddname option in a parameter list, or the ddname specified in a TO(ddname) option of an ICETOOL operator, as the default ddname for an OUTFIL statement without a FNAMES or FILES option.

z/OS V1R5 and Release 14 - PTFs (April, 2006)

INCLUDE and OMIT Enhancements

COND now allows you to compare date fields in various formats to past and future dates (relative to the date of your DFSORT run) using new DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r, DATEnP-r, Y'DATEn'+r and Y'DATEn'-r constants. &DATEn+r, &DATEn-r, &DATEn(c)+r, &DATEn(c)-r, &DATEnP+r and &DATEnP-r can be used as aliases for DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r and DATEnP-r, respectively.

COND now allows you to test a field for numerics (field,EQ,NUM) or non-numerics (field,NE,NUM) in character (FS), zoned decimal (ZD) or packed decimal (PD) format.

COND now allows you to use &DATEn, &DATEn(c) and &DATEnP as aliases for DATEn, DATEn(c) and DATEnP, respectively.

INREC and OUTREC Enhancements

PARSE and IFTHEN PARSE are new options that allow you to extract variable position/length fields into fixed-length parsed fields defined as %nn fields. PARSE gives you powerful new capabilities for handling variable fields such as delimited fields, comma separated values (CSV), tab separated values, blank separated values, keyword separated fields, null-terminated strings, and many other types.

You can use various PARSE options to define the rules for extracting variable fields into up to one hundred %nn fixed-length parsed fields (%00-%99), and then use these %nn fields where you can use p,m fields in BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and FIELDS. You can edit, convert, justify, squeeze, translate, and do arithmetic with %nn fields.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and FIELDS now allow you to use a new JFY option to left-justify or right-justify the data in a field. For a left-justified field, leading blanks are removed and the characters from the first nonblank to the last nonblank are shifted left, with blanks inserted on the right if needed. For a right-justified field, trailing blanks are removed and the characters from the last nonblank to the first nonblank are shifted right, with blanks inserted on the left if needed.

Optionally for JFY, specific leading and trailing characters can be changed to blanks before justification begins, a leading string can be inserted, a trailing string can be inserted, and the output length can be changed.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and FIELDS now allow you to use a new SQZ option to left-squeeze or right-squeeze the data in a field. For a left-squeezed field, all blanks are removed and the characters from the first nonblank to the last nonblank are shifted left, with blanks inserted on the right if needed. For a right-squeezed field, all blanks are removed and the characters from the last nonblank to the first nonblank are shifted right, with blanks inserted on the left if needed.

Optionally for SQZ, specific characters can be changed to blanks before squeezing begins, a leading string can be inserted, a trailing string can be inserted, a string can be inserted wherever a group of blanks is removed between the first nonblank and the last nonblank, blanks can be kept as is between paired apostrophes or paired quotes, and the output length can be changed.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and FIELDS now allow you to insert past and future dates (relative to the date of your DFSORT run) into your records in various forms using new DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r and DATEnP-r constants. &DATEn+r, &DATEn-r, &DATEn(c)+r, &DATEn(c)-r, &DATEnP+r and &DATEnP-r can be used as aliases for DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r and DATEnP-r, respectively.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and FIELDS now allow you to use FL format to convert 4-byte or 8-byte hexadecimal floating-point values to integer values.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and FIELDS now allow you to use new TO=PDF and TO=PDC options to convert numeric values to PD values with F or C for the positive sign, respectively. The TO=PDC option is equivalent to the existing TO=PD option.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and FIELDS now allow you to use &DATEn, &DATEn(c), &DATEnP, &YDDD=(abc), &YDDDNS=(ab), &TIMEn, &TIMEn(c) and &TIMEnP as aliases for DATEn, DATEn(c), DATEnP, YDDD=(abc), YDDDNS=(ab), TIMEn, TIMEn(c) and TIMEnP, respectively.

IFTHEN WHEN now allows you to compare date fields in various formats to past and future dates (relative to the date of your DFSORT run) using new DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r, DATEnP-r, Y'DATEn'+r and Y'DATEn'-r constants. &DATEn+r, &DATEn-r, &DATEn(c)+r, &DATEn(c)-r, &DATEnP+r and

&DATEnP-r can be used as aliases for DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r and DATEnP-r, respectively.

IFTHEN WHEN now allows you to test a field for numerics (field,EQ,NUM) or non-numerics (field,NE,NUM) in character (FS), zoned decimal (ZD) or packed decimal (PD) format.

IFTHEN WHEN now allows you to use &DATEn, &DATEn(c) and &DATEnP as aliases for DATEn, DATEn(c) and DATEnP, respectively.

OUTFIL Enhancements

BLKCCH1 is a new report option that allows you to avoid forcing a page eject at the start of the report header; the ANSI carriage control character of '1' (page eject) in the first line of the report header (HEADER1) is replaced with a blank.

BLKCCH2 is a new report option that allows you to avoid forcing a page eject at the start of the first page header; the ANSI carriage control character of '1' (page eject) in the first line of the first page header (HEADER2) is replaced with a blank.

BLKCCT1 is a new report option that allows you to avoid forcing a page eject at the start of the report trailer; the ANSI carriage control character of '1' (page eject) in the first line of the report trailer (TRAILER1) is replaced with a blank.

SPLITIR is a new option that allows you to write contiguous groups of records in one rotation among multiple output data sets. A specified number of records is written to each output data set and extra records are written to the last output data set.

PARSE and IFTHEN PARSE are new options that allow you to extract variable position/length fields into fixed-length parsed fields defined as %nn fields. PARSE gives you powerful new capabilities for handling variable fields such as delimited fields, comma separated values (CSV), tab separated values, blank separated values, keyword separated fields, null-terminated strings, and many other types.

You can use various PARSE options to define the rules for extracting variable fields into up to one hundred %nn fixed-length parsed fields (%00-%99), and then use these %nn fields where you can use p,m fields in BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and OUTREC. You can edit, convert, justify, squeeze, translate, and do arithmetic with %nn fields.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and OUTREC now allow you to use a new JFY option to left-justify or right-justify the data in a field. For a left-justified field, leading blanks are removed and the characters from the first nonblank to the last nonblank are shifted left, with blanks inserted on the right if needed. For a right-justified field, trailing blanks are removed and the characters from the last nonblank to the first nonblank are shifted right, with blanks inserted on the left if needed.

Optionally for JFY, specific leading and trailing characters can be changed to blanks before justification begins, a leading string can be inserted, a trailing string can be inserted, and the output length can be changed.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and OUTREC now allow you to use a new SQZ option to left-squeeze or right-squeeze the data in a field. For a left-squeezed field, all blanks are removed and the characters from the first nonblank to the last nonblank are shifted left, with blanks inserted on the right if needed. For a right-squeezed field, all blanks are removed and the characters from the last nonblank to the first nonblank are shifted right, with blanks inserted on the left if needed.

Optionally for SQZ, specific characters can be changed to blanks before squeezing begins, a leading string can be inserted, a trailing string can be inserted, a string can be inserted wherever a group of blanks is removed between the first nonblank and the last nonblank, blanks can be kept as is between paired apostrophes or paired quotes, and the output length can be changed.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and OUTREC now allow you to insert past and future dates (relative to the date of your DFSORT run) into your records in various forms using new DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r and DATEnP-r constants. &DATEn+r, &DATEn-r, &DATEn(c)+r, &DATEn(c)-r, &DATEnP+r and &DATEnP-r can be used as aliases for DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r and DATEnP-r, respectively.

TRAILERx, BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and OUTREC now allow you to use FL format to convert 4-byte or 8-byte hexadecimal floating-point values to integer values.

TRAILERx, HEADERx, BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and OUTREC now allow you to use new TO=PDF and TO=PDC options to convert numeric values to PD values with F or C for the positive sign, respectively. The TO=PDC option is equivalent to the existing TO=PD option.

BUILD, OVERLAY, IFTHEN BUILD, IFTHEN OVERLAY and OUTREC now allow you to use &DATEn, &DATEn(c), &DATEnP, &YDDD=(abc), &YDDDNS=(ab), &TIMEn, &TIMEn(c) and &TIMEnP as aliases for DATEn, DATEn(c), DATEnP, YDDD=(abc), YDDDNS=(ab), TIMEn, TIMEn(c) and TIMEnP, respectively.

INCLUDE, OMIT and IFTHEN WHEN now allow you to compare date fields in various formats to past and future dates (relative to the date of your DFSORT run) using new DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r, DATEnP-r, Y'DATEn'+r and Y'DATEn'-r constants. &DATEn+r, &DATEn-r, &DATEn(c)+r, &DATEn(c)-r, &DATEnP+r and &DATEnP-r can be used as aliases for DATEn+r, DATEn-r, DATEn(c)+r, DATEn(c)-r, DATEnP+r and DATEnP-r, respectively.

INCLUDE, OMIT and IFTHEN WHEN now allow you to test a field for numerics (field,EQ,NUM) or non-numerics (field,NE,NUM) in character (FS), zoned decimal (ZD) or packed decimal (PD) format.

INCLUDE, OMIT and IFTHEN WHEN now allow you to use &DATEn, &DATEn(c) and &DATEnP as aliases for DATEn, DATEn(c) and DATEnP, respectively.

Symbol Enhancements

A symbol can now be used for a %nn parsed field. For example, if Account,%01 is defined in SYMNAMES, Account can be used for %01. A symbol for %nn can be used in DFSORT control statements where %nn can be used. A symbol for %nn results in substitution of %nn.

A symbol can now be used for an output column. For example, if Start_address,18 is defined in SYMNAMES, Start_address: can be used for 18:. symbol: can be used in DFSORT control statements where c: can be used. A symbol for p or p,m or p,m,f results in substitution of p: for symbol: (output column).

A symbol can now be used for a new system symbol string constant. symbol,S'string' can be used to define a string containing any combination of EBCDIC characters and system symbols you want to use to form a character string. For example, if whererun,S'&JOBNAME. on &SYSPLEX' is defined in SYMNAMES, whererun can be used for the resulting constant. You can use dynamic system symbols such as &JOBNAME, &DAY, and so on, system-defined static system symbols such as &SYSNAME, &SYSPLEX, and so on, and installation-defined static system symbols specified by your installation in an IEASYMxx member of SYS1.PARMLIB.

A symbol for a system symbol string can be used in DFSORT and ICETOOL control statements where a symbol for a character string can be used. DFSORT will replace each system symbol in S'string' with its substitution text to create a character string in the format C'new_string'.

ICETOOL Enhancements

DISPLAY now allows you to use FL format to convert 4-byte or 8-byte hexadecimal floating-point values to integer values.

DISPLAY and OCCUR now allow you to use a new TBETWEEN(n) option to specify the number of blanks between title elements (title, page number, date, time).

SELECT and SPLICE now allow you to use an INREC statement to reformat your records before they are selected or spliced. All of the operands of the INREC statement (PARSE, BUILD, OVERLAY, IFTHEN, IFOUTLEN and FIELDS) are now available with SELECT and SPLICE.

SORT and MERGE Enhancements

The maximum length for a PD or ZD sort or merge field has been raised to 256.

SUM Enhancements

The maximum position for the end of a sum field has been raised to 32752.

Other Enhancements

DFSORT supports large physical sequential data sets for input, output and work data sets.

DSA can now be specified as a run-time option. This allows you to adjust the maximum amount of storage available to DFSORT for dynamic storage adjustment of individual Blockset sort applications when SIZE/MAINSIZE=MAX is in effect.

DFSORT now accepts and ignores zero values in the starting and ending address of the RECORD statement image in the 24-Bit Parameter List. You can set these addresses to zero if you don't want to pass a control statement to DFSORT using the third and fourth words of the parameter list .

z/OS V1R5 and Release 14 - PTFs (December, 2004)

ICETOOL Enhancements

ICETOOL's DISPLAY, OCCUR, RANGE, SELECT, SPLICE, STATS, UNIQUE and VERIFY operators now allow you to use larger numeric values for ON and BREAK fields. PD, ZD and FS fields can now be up to 31 digits (or more in some cases). BI and FI fields can now be up to 8 bytes (or more in some cases).

ICETOOL's RANGE operator now allows you to use larger decimal values for the HIGHER(n), LOWER(n), EQUAL(n) and NOTEQUAL(n) options. These values can now be up to 31 digits.

ICETOOL's DISPLAY, OCCUR, RANGE, SELECT, SPLICE, STATS and UNIQUE operators now allow you to use new UFF (unsigned free form) and SFF (signed free form) formats for ON and BREAK fields. UFF extracts a

positive numeric value from a free form field (for example, '\$1234.56' is treated as +123456). SFF extracts a positive or negative value from a free form field (e.g., '(1,234.56)' is treated as -123456).

ICETOOL's DISPLAY and OCCUR operators now allow you to use new DC1-DC3 (TOD date), DE1-DE3 (ETOD date), TC1-TC4 (TOD time) and TE1-TE4 (ETOD time) formats for ON and BREAK fields. These new formats produce meaningful representations of TOD and ETOD date and time values.

ICETOOL's DISPLAY and OCCUR operators now allow you to specify multiline headings for the columns of your reports. You can specify one, two or three line headings with the HEADER option.

ICETOOL's DISPLAY and OCCUR operators now allow you to use new edit masks G1-G6 to display numeric values with 4 decimal digits in various ways.

ICETOOL's DISPLAY and OCCUR operators now allow you to use new YDDD(abc) and YDDDNS(abc) options to insert the year (yyyy) and day of the year (ddd) of your ICETOOL run into your titles in various forms.

ICETOOL's COUNT operator now allows you to use a new RC4 option to set RC=4 (instead of RC=12) or RC=0 based on the count of records in a data set.

ICETOOL's SPLICE operator now allows you to use a new KEEPBASE option to keep the base records as well as the spliced records.

ICETOOL's SPLICE operator now allows you to use new VLENMAX or VLENOVLY options to set the length of spliced records to the maximum of the base or overlay record length, or to the overlay record length.

OUTFIL Enhancements

OUTFIL now allows you to reformat records in one of the following three ways using unedited, edited, or converted input fields and a variety of constants:

- BUILD or OUTREC: The existing OUTREC parameter, or its new alias of BUILD, allows you to reformat each record by specifying all of its items one by one. BUILD or OUTREC gives you complete control over the items you want in your reformatted OUTFIL records and the order in which they appear. You can delete, rearrange and insert fields and constants.
- OVERLAY: The new OVERLAY parameter allows you to reformat each record by specifying just the items that overlay specific columns. Overlay lets you change specific existing columns without affecting the entire record.
- IFTHEN clauses: The new IFTHEN clauses allow you to reformat different records in different ways by specifying how BUILD or OVERLAY items are applied to records that meet given criteria. IFTHEN clauses let you use sophisticated conditional logic to choose how different record types are reformatted.

OUTFIL OUTREC, as well as BUILD, OVERLAY and IFTHEN, now allows you to use larger numeric values for fields and decimal constants to be edited, converted or used in arithmetic expressions. PD, ZD and FS fields, and decimal constants, can now be up to 31 digits. BI and FI fields can now be up to 8 bytes.

OUTFIL OUTREC, as well as BUILD, OVERLAY and IFTHEN, now allows you to use new UFF (unsigned free form) and SFF (signed free form) formats for fields to be edited, converted or used in arithmetic expressions.

OUTFIL OUTREC, as well as BUILD, OVERLAY and IFTHEN, now allows you to use new DC1-DC3 (TOD date), DE1-DE3 (ETOD date), TC1-TC4 (TOD time) and TE1-TE4 (ETOD time) formats for fields to be edited, converted or used in arithmetic expressions. These new formats produce meaningful representations of TOD and ETOD date and time values.

OUTFIL OUTREC , as well as BUILD, OVERLAY and IFTHEN, now allows you to use a set field in the CHANGE option (e.g., 1,2,CHANGE=(4,C'FY',C'0001',C'VV',21,4)).

OUTFIL OUTREC, as well as BUILD, OVERLAY and IFTHEN, now allows you to restart the sequence number when the binary value of a specified field changes (e.g., SEQNUM,5,ZD,RESTART=(11,4)).

OU2FIL OUTREC, as well as BUILD, OVERLAY and IFTHEN, now allows you to use DATE, DATE=(abcd), DATENS=(abc), YDDD=(abc), YDDDNS=(ab), TIME, TIME=(abc) and TIMENS=(ab) options to insert the date and time of your DFSORT run into your records in various forms.

OUTFIL OUTREC, as well as BUILD, OVERLAY and IFTHEN, now allows you to use new TO=ZDC and TO=ZDF options to convert numeric values to ZD values with C or F for the positive sign, respectively. The TO=ZDF option is equivalent to the existing TO=ZD option.

OUTFIL INCLUDE and OMIT now allow you to use larger FS values for compare fields. These values can now be up to 32 digits.

OUTFIL INCLUDE and OMIT now allow you to use larger decimal constants for comparison to BI and FI fields. Decimal constants can now be up to +18446744073709551615 for comparison to BI fields. Decimal constants can now be between -9223372036854775808 and +9223372036854775807 for comparison to FI fields.

OUTFIL INCLUDE and OMIT now allow you to use new UFF (unsigned free form) and SFF (signed free form) formats for compare fields. A UFF or SFF field can be compared to a UFF, SFF, FS, CSL or CST field or to a decimal constant.

OUTFIL TRAILERx now allows you to use larger numeric values for statistical fields (total, maximum, minimum, average). PD, ZD and FS fields, and decimal constants, can now be up to 31 digits. BI and FI fields can now be up to 8 bytes.

OUTFIL TRAILERx now allows you to use new UFF (unsigned free form) and SFF (signed free form) formats for statistical fields (total, maximum, minimum, average).

OUTFIL TRAILERx now allows you to use TO=fo and fo (to) options to convert statistical fields (total, maximum, minimum, average) and counts to BI, FI, PD, ZD, ZDC, ZDF or FS output values.

OUTFIL TRAILERx now allows you to use COUNT+n=(edit), COUNT+n=(to), COUNT-n=(edit) and COUNT-n=(to) to add or subtract n from a count to be edited or converted (e.g., COUNT+1=(TO=ZD)).

OUTFIL HEADERx and TRAILERx now allow you to insert hexadecimal strings (X'yy...yy' or nX'yy...yy') in your headers and trailers.

OUTFIL HEADERx and TRAILERx now allow you to use new YDDD=(abc) and YDDDNS=(abc) options to insert the year (yyyy) and day of the year (ddd) of your DFSORT run in your headers and trailers.

OUTFIL HEADERx and TRAILERx now allow you to use PAGE=(edit) and PAGE=(to) to edit or convert the page number (e.g., PAGE=(M11,LENGTH=3)).

INREC and OUTREC Enhancements

INREC and OUTREC now allow you to reformat records in one of the following three ways using unedited, edited, or converted input fields and a variety of constants:

- i BUILD or FIELDS: The existing FIELDS parameter, or its new alias of BUILD, allows you to reformat each record by specifying all of its items one by one. BUILD or FIELDS gives you complete control over the items

you want in your reformatted INREC or OUTREC records and the order in which they appear. You can delete, rearrange and insert fields and constants.

- **OVERLAY:** The new OVERLAY parameter allows you to reformat each record by specifying just the items that overlay specific columns. Overlay lets you change specific existing columns without affecting the entire record.
- **IFTHEN clauses:** The new IFTHEN clauses allow you to reformat different records in different ways by specifying how BUILD or OVERLAY items are applied to records that meet given criteria. IFTHEN clauses let you use sophisticated conditional logic to choose how different record types are reformatted.

INREC and OUTREC now allow you to use larger numeric values for fields and decimal constants to be edited, converted or used in arithmetic expressions. PD, ZD and FS fields, and decimal constants, can now be up to 31 digits. BI and FI fields can now be up to 8 bytes.

INREC and OUTREC now allow you to use new UFF (unsigned free form) and SFF (signed free form) formats for fields to be edited, converted or used in arithmetic expressions.

INREC and OUTREC now allow you to use new DC1-DC3 (TOD date), DE1-DE3 (ETOD date), TC1-TC4 (TOD time) and TE1-TE4 (ETOD time) formats for fields to be edited, converted or used in arithmetic expressions. These new formats produce meaningful representations of TOD and ETOD date and time values.

INREC and OUTREC now allow you to use a set field in the CHANGE option.

INREC and OUTREC now allow you to restart the sequence number when the binary value of a specified field changes.

INREC and OUTREC now allow you to use new TO=ZDC and TO=ZDF options to convert numeric values to ZD values with C or F for the positive sign, respectively. The TO=ZDF option is equivalent to the existing TO=ZD option.

INREC and OUTREC now allow you to use new DATE, DATE=(abcd), DATENS=(abc), YDDD=(abc), YDDDNS=(ab), TIME, TIME=(abc) and TIMENS=(ab) options to insert the date and time of your DFSORT run into your records in various forms.

SORT and MERGE Enhancements

SORT and MERGE now allow you to use larger FS values for control fields. These values can now be up to 32 digits.

SORT and MERGE now allow you to use new UFF (unsigned free form) and SFF (signed free form) formats for control fields.

SORT and MERGE now allow you to use other characters besides '+' (plus) as a positive sign in control field values for CSL, CST, ASL and AST formats. A '-' (minus) sign is treated as negative and any other sign (e.g., blank) is treated as positive.

INCLUDE and OMIT Enhancements

INCLUDE and OMIT now allow you to use larger FS values for compare fields. These values can now be up to 32 digits.

INCLUDE and OMIT now allow you to use larger decimal constants for comparison to BI and FI fields. Decimal constants can now be up to +18446744073709551615 for comparison to BI fields. Decimal constants can now be between -9223372036854775808 and +9223372036854775807 for comparison to FI fields.

INCLUDE and OMIT now allow you to use new UFF (unsigned free form) and SFF (signed free form) formats for compare fields. A UFF or SFF field can be compared to a UFF, SFF, FS, CSL or CST field or to a decimal constant.

SUM Enhancement

SUM now allows you to use larger ZD values for sum fields. These values can now be up to 31 digits.

Other Enhancements

For sort and copy applications, with concatenated variable-length input data sets for SORTIN, DFSORT now uses the largest LRECL it finds in the concatenation.

DFSORT now allows control statements in SYSIN and SORTCNTL to contain labels up to 70 characters, and allows any character in the label. DFSORT now ignores statements with a label followed only by blanks.

z/OS V1R5 - Base (March, 2004)

Improvements in Performance

Memory object sorting is a new DFSORT capability that uses a memory object on 64-bit real architecture to improve the performance of sort applications. A memory object is a data area in virtual storage that is allocated above the bar and backed by central storage. With memory object sorting, a memory object can be used exclusively, or along with disk space, for temporary storage of records. Memory object sorting can reduce I/O processing, elapsed time, EXCPs, and channel usage. When a memory object is used, Hiperspace and data space are not needed.

Installation and Run-Time Options

A new MOSIZE installation (ICEMAC) and run-time option allows you to specify the maximum size of a memory object to be used for memory object sorting. The existing EXPMAX, EXPOLD and EXPRES installation (ICEMAC) options now apply to memory object sorting as well as Hipersorting.

A new NULLOUT installation (ICEMAC) and run-time option allows you to specify what you want DFSORT to do when there are no records for the SORTOUT data set. This gives you control over the action (continue or terminate), type of message (informational or error) and return code (0, 4 or 16) for a SORTOUT data set with no records.

A new NULLOFL installation (ICEMAC) and OUTFIL run-time option allows you to specify what you want DFSORT to do when there are no records for an OUTFIL data set. This gives you control over the action (con-

tinue or terminate), type of message (informational or error) and return code (0, 4 or 16) for an OUTFIL data set with no data records.

The IBM-supplied default for ICEMAC option COBEXIT has been changed from COB1 to COB2.

The IBM-supplied default for ICEMAC option DSA has been changed from 32MB to 64MB.

The IBM-supplied default for ICEMAC option TMAXLIM has been changed from 4MB to 6MB.

The IBM-supplied default for ICEMAC option ZDPRINT has been changed from NO to YES.

Release 14 - PTFs (February, 2003)

ICETOOL Enhancements

A new SPLICE operator helps you to perform various file "join" and "match" operations. SPLICE allows you to create output records in a variety of ways by splicing together fields from records that have the same key, but different information. For example, for the same userid, you could create output records consisting of the division and department from one type of record (originating in input file1) with the account number and phone number from another type of record (originating in input file2). Spliced records can be created by combining the first duplicate and:

- one or more fields from the last duplicate, or
- one or more fields from each subsequent duplicate, or
- one field from each subsequent duplicate.

Non-duplicate records can be deleted or kept.

The USING(yyyy) option can now be used with ICETOOL's SELECT operator to process DFSORT control statements like INCLUDE, OMIT and OUTFIL for a SELECT operation.

New LISTSDB and LISTNOSDB options for ICETOOL's DEFAULTS, DISPLAY and OCCUR operators allow you to control the use of system-determined optimum blocksize for LIST data sets.

OUTFIL Enhancements

New SAMPLE=n and SAMPLE=(n,m) options of OUTFIL allow you to sample records in a variety of ways.

A new REPEAT=n option of OUTFIL allows you to write each output record multiple times.

A new SPLITBY=n option of OUTFIL allows you to write groups of records in rotation among multiple output data sets.

OUTFIL OUTREC now allows you to insert decimal constants (+n and -n) in your records as BI, FI, PD, ZD, FS or edited CH values.

OUTFIL OUTREC now allows you to combine fields (p,m,f), decimal constants (+n and -n), operators (MIN, MAX, MUL, DIV, MOD, ADD, SUB) and parentheses to form arithmetic expressions, and place the results in your records as BI, FI, PD, ZD, FS or edited CH values.

A new DATE4 option of OUTFIL OUTREC allows you to insert a timestamp for your DFSORT run in the form 'yyyy-mm-dd-hh.mm.ss' into your records.

The maximum length for an SS field used with OUTFIL INCLUDE and OUTFIL OMIT has been raised to 32752.

A new DATE4 option of OUTFIL INCLUDE and OUTFIL OMIT allows you to compare fields to a timestamp for your DFSORT run in the form 'yyyy-mm-dd-hh.mm.ss' or to a portion of that timestamp truncated on the right.

A PD0 field can now be compared to a hexadecimal constant or to another PD0 field for OUTFIL INCLUDE and OMIT.

OUTFILE can now be used as an alias for OUTFIL.

INREC and OUTREC Enhancements

INREC and OUTREC now allow you to insert decimal constants (+n and -n) in your records as BI, FI, PD, ZD, FS or edited CH values.

INREC and OUTREC now allow you to combine fields (p,m,f), decimal constants (+n and -n), operators (MIN, MAX, MUL, DIV, MOD, ADD, SUB) and parentheses to form arithmetic expressions, and place the results in your records as BI, FI, PD, ZD, FS or edited CH values.

A new DATE4 option of INREC and OUTREC allows you to insert a timestamp for your DFSORT run in the form 'yyyy-mm-dd-hh.mm.ss' into your records.

INCLUDE and OMIT Enhancements

The maximum length for an SS field used with INCLUDE and OMIT has been raised to 32752.

A new DATE4 option of INCLUDE and OMIT allows you to compare fields to a timestamp for your DFSORT run in the form 'yyyy-mm-dd-hh.mm.ss' or to a portion of that timestamp truncated on the right.

A PD0 field can now be compared to a hexadecimal constant or to another PD0 field for INCLUDE and OMIT.

FORMAT=f can now be used with mixed p,m and p,m,f fields in the COND operand for INCLUDE and OMIT. f from FORMAT=f will be used for p,m fields but not for p,m,f fields.

SORT, MERGE and SUM Enhancement

FORMAT=f can now be used with mixed p,m and p,m,f fields in the FIELDS operand for SORT, MERGE and SUM. f from FORMAT=f will be used for p,m fields but not for p,m,f fields.

Other Enhancements

Enhancements to DFSORT's control statement continuation rules allow you to continue a line that breaks at column 71 anywhere in columns 2 to 16 of the next line.

When PARMDDN=ddname is specified at installation-time, DFSORT will now use a //DFSPARM DD data set if a //ddname DD data set is not present. When PARMDDN=DFSPARM is specified or defaulted at installation-time, DFSORT will continue to use a //\$ORTPARM DD data set if a //DFSPARM DD data set is not present.

Release 14 - PTFs (March, 2002)

z900

DFSORT can now exploit the z900's 64-bit real architecture by backing storage and data spaces in real storage above 2 gigabytes, and by using central storage instead of expanded storage for Hipersorting.

Multiple Hiperspaces

DFSORT can now use multiple Hiperspaces for external storage requirements, increasing DFSORT's ability to use Hipersorting for large sort applications when sufficient system resources are available.

Managed Tapes

DFSORT can now automatically obtain accurate input file size information for tapes managed by tape management systems that supply a tape exit routine (ICETPEX). This can result in improved sort performance and more accurate dynamic work space allocation.

DFSORT can now automatically obtain input and output attributes such as RECFM, LRECL and BLKSIZE, for tapes managed by tape management systems that supply a tape exit routine (ICETPEX).

ICETOOL Enhancements

New EMPTY, NOTEMPTY, HIGHER(n), LOWER(n), EQUAL(n) and NOTEQUAL(n) operands of ICETOOL's COUNT operator allow you to set RC=12 or RC=0 based on the count of records in a data set.

A new FIRSTDUP option of ICETOOL's SELECT operator allows you to keep just the first record of each set of duplicates.

A new LASTDUP option of ICETOOL's SELECT operator allows you to keep just the last record of each set of duplicates.

The DISCARD(savedd) option of ICETOOL's SELECT operator can now be specified without the TO(outdd) operand.

The maximum length for a CH field used with ICETOOL has been raised to 1500.

The maximum length for a HEX field used with ICETOOL has been raised to 1000.

New DT1, DT2 and DT3 formats of ICETOOL's DISPLAY and OCCUR operators allow you to produce meaningful representations of SMF date values.

New TM1, TM2, TM3 and TM4 formats of ICETOOL's DISPLAY and OCCUR operators allow you to produce meaningful representations of SMF time values.

A new INDENT(n) option of ICETOOL's DISPLAY and OCCUR operators allows you to specify the number of blanks for indentation.

A new BETWEEN(n) option of ICETOOL's DISPLAY and OCCUR operators allows you to specify the number of blanks between columns.

A new STATLEFT option of ICETOOL's DISPLAY operator allows you to print the headings for statistics to the left of the first column of data.

The use of formatting items has been extended to the record length, record number and break field of ICETOOL's DISPLAY operator, and to the data field, record length and value count of ICETOOL's OCCUR operator.

Additional formatting items are now provided as follows: E'pattern' (use a specified pattern for numeric digits such as phone numbers, social security numbers, dates, and so on), /D (divide by 10), /C (divide by 100), /DK (divide by 10000), /CK (divide by 100000), LZ (leading zeros for mask), NOST (no statistics), and Ndd (dd digits).

A new DATENS(abc) option of ICETOOL's DISPLAY and OCCUR operators allows you to insert the date of your ICETOOL run into your titles without a separator. Separators are included for existing options DATE and DATE(abcd).

A new TIMENS(ab) option of ICETOOL's DISPLAY and OCCUR operators allows you to insert the time of your ICETOOL run into your titles without a separator. Separators are included for existing options TIME and TIME(abc).

OUTFIL Enhancements

A new FTOV option of OUTFIL allows you to convert fixed-length input records (for example, FB) to variable-length output records (for example, VB). A new VTOF option of OUTFIL (an alias for CONVERT) allows you to convert variable-length input records (for example, VB) to fixed-length output records (for example, FB).

A new VLTRIM=byte option of OUTFIL allows you to remove trailing bytes such as blanks, binary zeros or asterisks, from variable-length output records.

A new REMOVECC option of OUTFIL allows you to remove the ANSI control characters from a report.

New DATEn, DATEn(c) and DATEnP options of OUTFIL OUTREC allow you to insert the date of your DFSORT run into your records in various forms.

New TIMEn, TIMEn(c) and TIMEnP options of OUTFIL OUTREC allow you to insert the time of your DFSORT run into your records in various forms.

New DT1, DT2 and DT3 formats of OUTFIL OUTREC allow you to produce meaningful representations of SMF date values.

New TM1, TM2, TM3 and TM4 formats of OUTFIL OUTREC allow you to produce meaningful representations of SMF time values.

A new TRAN=LTOU option of OUTFIL OUTREC allows you to change lowercase EBCDIC letters anywhere in a field to uppercase EBCDIC letters.

A new TRAN=UTOL option of OUTFIL OUTREC allows you to change uppercase EBCDIC letters anywhere in a field to lowercase EBCDIC letters.

A new TRAN=ALTSEQ option of OUTFIL OUTREC allows you to change characters anywhere in a field according to the ALTSEQ translation table in effect.

A new M26 edit mask of OUTFIL OUTREC allows you to edit numeric values according to the pattern ST...T (S is the sign and T is a significant digit).

The maximum position for the end of a field used with `OUTFIL INCLUDE` and `OMIT` has been raised to 32752.

New `DATEn`, `DATEn(c)` and `DATEnP` options of `OUTFIL INCLUDE` and `OMIT` allow you to compare fields with various formats to the date of your DFSORT run.

A positive decimal number (`n` or `+n`) can now be compared to a binary (BI) field for `OUTFIL INCLUDE` and `OMIT`.

A new `DATENS=(abc)` option of `OUTFIL HEADERx` and `TRAILERx` allows you to insert the date of your DFSORT run into your headers and trailers without a separator. Separators are included for existing options `DATE` and `DATE=(abcd)`.

A new `TIMENS=(ab)` option of `OUTFIL HEADERx` and `TRAILERx` allows you to insert the time of your DFSORT run into your headers and trailers without a separator. Separators are included for existing options `TIME` and `TIME=(abc)`.

A new `SKIP=0L` option of `OUTFIL SECTIONS` allows you to suppress blank lines between sections on the same page.

INREC and OUTREC Enhancements

New `DATEn`, `DATEn(c)` and `DATEnP` options of `INREC` and `OUTREC` allow you to insert the date of your DFSORT run into your records in various forms.

New `TIMEn`, `TIMEn(c)` and `TIMEnP` options of `INREC` and `OUTREC` allow you to insert the time of your DFSORT run into your records in various forms.

New `DT1`, `DT2` and `DT3` formats of `INREC` and `OUTREC` allow you to produce meaningful representations of SMF date values.

New `TM1`, `TM2`, `TM3` and `TM4` formats of `INREC` and `OUTREC` allow you to produce meaningful representations of SMF time values.

A new `TRAN=LTOU` option of `INREC` and `OUTREC` allows you to change lowercase EBCDIC letters anywhere in a field to uppercase EBCDIC letters.

A new `TRAN=UTOL` option of `INREC` and `OUTREC` allows you to change uppercase EBCDIC letters anywhere in a field to lowercase EBCDIC letters.

A new `TRAN=ALTSEQ` option of `INREC` and `OUTREC` allows you to change characters anywhere in a field according to the ALTSEQ translation table in effect.

A new `M26` edit mask of `INREC` and `OUTREC` allows you to edit numeric values according to the pattern `ST...T` (S is the sign and T is a significant digit).

SUM Enhancement

The `VLSHRT` option can now be used to process short SORT and MERGE control fields, and `INCLUDE` and `OMIT` compare fields, when a `SUM` statement is specified. Records with short SUM summary fields are left unsummed.

SORT and MERGE Enhancements

The maximum position for the end of a field used with SORT and MERGE has been raised to 32752.

The maximum length for an AQ or AC field used with SORT and MERGE has been raised to 4092.

INCLUDE and OMIT Enhancements

New DATEn, DATEn(c) and DATEnP options of INCLUDE and OMIT allow you to compare fields with various formats to the date of your DFSORT run.

The maximum position for the end of a field used with INCLUDE and OMIT has been raised to 32752.

A positive decimal number (n or +n) can now be compared to a binary (BI) field for INCLUDE and OMIT.

Other Enhancements

The IBM-supplied default for ICEMAC option DYNSPC has been changed from 32 megabytes to 256 megabytes. This allows DFSORT to sort more data by default with dynamically allocated work space when the input file size is unknown.

DYNSPC can now be specified as a run-time option. This allows you to adjust the dynamically allocated work space for individual sort applications for which the input file size is unknown.

DFSORT now issues message ICE118I, and sets on bit 3 of SMF type-16 record field ICEFSZFL, when the input file size is unknown for a sort application. This identifies sort applications which may benefit from a FILSZ=En or DYNSPC=n value.

DFSORT copy can now be used when ICEGENER is called by a program that uses an alternate SYSIN ddname with DUMMY. This can result in improved performance for RACF's IRRUT200 utility when ICEGENER is installed as a replacement for IEBGENER.

When all of the input for variable-length records is supplied through an E15 or E32 exit, and RECFM is not specified for the output data set, DFSORT now sets the output RECFM as blocked rather than unblocked, if that allows the use of the system-determined optimum block size for output.

A hyphen (-) can now be used as the second or subsequent character in a DFSORT symbol name.

DFSORT now issues message ICE042A and terminates if the QNAME option is specified on an output DD statement.

Release 14 - PTFs (July, 2000)

New Device Support

DFSORT now supports FICON channels.

DFSORT now recognizes the IBM 2105 control unit and provides optimizations for it appropriate for cached control units, such as cache fast write for work data sets.

DFSORT now treats work data sets on emulated 3390 model 9 devices as desirable choices with respect to performance. DFSORT continues to treat work data sets on real 3390 model 9 devices as undesirable choices with respect to performance.

Larger Tape Block Sizes with OS/390 R10

DFSORT can now use tape data sets with block sizes greater than 32760 bytes for input and output, providing improved performance and tape utilization.

DFSORT can now select system-determined optimum block sizes greater than 32760 bytes for tape output data sets, if allowed by the BLKSZLIM value in effect. Installation and run-time options SDB=INPUT (the new IBM-supplied default), SDB=LARGE (new), SDB=YES (or its alias SDB=SMALL) and SDB=NO allow you to control DFSORT's use of system-determined block sizes, including block sizes greater than 32760 bytes for tape output data sets.

DFSORT's ICEGENER, like IEBGENER, will use SDB=value parameters you specify.

DFSMSrmm-Managed Tapes

DFSORT can now automatically obtain accurate input file size information for DFSMSrmm-managed tapes. This can result in improved sort performance and more accurate dynamic workspace allocation.

DFSORT can now automatically obtain input and output attributes such as RECFM, LRECL and BLKSIZE for DFSMSrmm-managed tapes.

VSAM Processing

DFSORT now supports the VSAM extended addressability function for extended format VSAM data sets, which provides the capability of VSAM data sets larger than four gigabytes.

DFSORT now allows an empty VSAM input data set to be accepted and processed as having zero records. New installation option VSAMEMT=YES/NO and run-time options VSAMEMT and NVSAMEMT allow you to control the processing of empty VSAM input data sets.

DFSORT now allows a VSAM output data set defined with REUSE to be processed as a NEW or MOD data set. New installation option RESET=YES/NO and run-time options RESET and NORESET allow you to control the processing of VSAM output data sets.

DFSORT now allows a VSAM data set defined with REUSE to be sorted inplace. New installation option VSAMIO=YES/NO and run-time options VSAMIO and NOVAMIO allow you to control the use of the same VSAM data set for input and output.

HFS Files

DFSORT now supports Hierarchical File System (HFS) files for input and output.

Productivity and Easier Migration from other Sort Products

By default, DFSORT now uses the L3, OUTREC or INREC record length as the SORTOUT LRECL when the SORTOUT LRECL is unavailable. New installation option SOLRF=YES/NO and run-time options SOLRF and NOSOLRF allow you to control the use of the reformatted record length for the SORTOUT LRECL.

DFSORT now provides enhanced capabilities for OUTFIL, including:

- Padding short records with blanks, by default, when converting variable-length records to fixed-length records.
- Creation of edited counts and subcounts in trailers.

DFSORT now provides enhanced capabilities for INREC and OUTREC that were previously only available with OUTFIL OUTREC, including:

- Hexadecimal display.
- Transformation of dates.
- Editing of numeric fields.
- Lookup and change.

DFSORT now provides enhanced capabilities for INREC, OUTREC and OUTFIL OUTREC, including:

- Conversion of numeric fields to other numeric formats.
- Editing of transformed dates.
- Conversion of transformed dates to numeric formats.
- Creation of sequence numbers.

DFSORT now allows INCLUDE/OMIT fields that extend beyond variable-length records (that is, short fields) to be compared as if they were padded with binary zeros. This allows both the short and non-short fields to be compared. New installation option VLSCMP=YES/NO and run-time options VLSCMP and NOVLSCMP allow you to control the use of binary zero padding for short control fields.

DFSORT now allows -0 and +0 to be treated as unsigned (that is, the same) for processing by DFSORT statements INCLUDE, INREC, MERGE, OMIT, OUTFIL, OUTREC and SORT. New installation option SZERO=YES/NO and run-time options SZERO and NOSZERO allow you to select signed or unsigned zero processing for these statements.

DFSORT now allows -0 and +0 to be treated as unsigned (that is, the same) for processing by ICETOOL operators DISPLAY, OCCUR, SELECT and UNIQUE. New option UZERO allows you to select unsigned zero processing for these operators.

DFSORT now always determines or assigns the record type (F or V) when the input is VSAM, or when an E15 or E32 exit supplies all of the input records.

DFSORT now allows an LRECL greater than 80 for SYSIN, SORTCNTL and DFSPARM; only the first 80 bytes of each control statement are used.

DFSORT now allows variable-length output records that are longer than the output data set LRECL to be truncated to the LRECL. New installation option VLLONG=YES/NO and run-time options VLLONG and NOVLLONG allow you to control the truncation of long variable-length output records.

DFSORT now accepts and processes the following as EXEC/DFSPARM PARM options: COBEXIT=COB1/COB2, NOVERIFY, NOVLSHRT, NZDPRINT, OVFL0=RC0/RC4/RC16, PAD=RC0/RC4/RC16, SPANINC=RC0/RC4/RC16, TRUNC=RC0/RC4/RC16, VERIFY, VLSHRT and ZDPRINT.

DFSORT now accepts and processes the following as OPTION statement options: EXITCK=STRONG/WEAK, WRKREL and WRKSEC.

DFSORT now accepts and processes Y2PAST=s/f as a SORT and MERGE statement option.

DFSORT now accepts, and processes or ignores as appropriate, many more aliases for DFSORT options.

DFSORT now accepts and processes the following as 24-Bit Parmlist control statement entry codes: X'10', X'16', X'18' and X'20' through X'29'.

New Generation of Year 2000 Features

New Y2T, Y2U, Y2V, Y2X, and Y2Y formats handle CH, ZD, and PD full dates like yymmdd, yyddd, yymm, yyq, mmdyy, dddyy, mmyy, and qyy, and their special indicators like zeros and nines, according to the specified fixed or sliding century window.

You can use these new formats in the SORT, MERGE, INCLUDE, OMIT, and OUTFIL statements to sort, merge, compare, and transform full dates.

Release 14 - Base (September, 1998)

Symbols for Fields and Constants

DFSORT now provides a simple and flexible method for using symbols in DFSORT and ICETOOL statements. You can define and use a symbol for any field or constant that is recognized in a DFSORT control statement or ICETOOL operator. This makes it easy to create and reuse collections of symbols (that is, mappings) representing information associated with various record layouts.

In addition, you can obtain and use collections of DFSORT symbols created specifically for records produced by other products (for example, RACE, DFSMSrmm and DCOLLECT) or by your site. These mappings can increase your productivity by automatically providing the positions, lengths and formats of the fields, and the values of the constants, associated with the particular records you are processing with DFSORT or ICETOOL.

Improvements in Performance, Capacity and Storage Usage

Blockset copy and merge applications can now use storage above 16 MB virtual, providing improved performance and virtual storage constraint relief.

Blockset copy and merge modules will now reside above 16 MB virtual, providing virtual storage constraint relief.

DFSORT can now handle a significantly larger number of INCLUDE and OMIT conditions.

DFSORT can now handle a significantly larger number of SUM fields.

The upper limit for the number of JCL and dynamically allocated work data sets that can be specified and used by DFSORT's Blockset technique has been raised from 100 to 255. The use of more work data sets increases the maximum amount of data DFSORT can process in a single sort application. Any valid ddname of the form

SORTWKdd or SORTWKd can now be used for DASD work data sets (for example, SORTWK01, SORTWKC3, SORTWK2, SORTWK#5, SORTWKA, SORTWKXY and so on).

The upper limit for the number of input data sets that can be specified and used for a Blockset merge application has been raised from 16 to 100. The use of more merge input data sets increases the maximum amount of data DFSORT can process in a single merge application.

Time-of-Day Option Controls

New time-of-day installation modules (ICETD1-4) allow different sets of installation defaults to be used, based on the day and time DFSORT applications run. Each environment installation module (ICEAM1-4) can enable one or more time-of-day installation modules. This capability allows new levels of control for installation defaults. For example, larger storage, hiperspace and data space limits could be used only for batch program-invoked DFSORT applications that run off-shift during the week, and all weekend.

Repackaging

The product has been repackaged to simplify installation and ordering:

- i The number of FMIDs has been reduced significantly.
- i The number of libraries for installation has been reduced significantly.
- i DFSORT R14 now supports a single installation of the product for both resident and nonresident features. This allows you to decide how to use DFSORT independent of the installation method, thus reducing the number of decisions you have to make at installation time.

OUTFIL Processing Enhancements

OUTFIL now supports creation of multiple output records from each input record, with or without intervening blank output records.

OUTFIL now supports processing of variable-length input records which are too short to contain all specified OUTFIL OUTREC fields. OUTFIL's new VLFILL=byte operand can be used to replace missing bytes in OUTFIL OUTREC fields with the specified fill byte so the filled fields can be processed.

ICETOOL Enhancement

A new DISCARD(savedd) operand of ICETOOL's SELECT operator allows you to save the records that are not selected, in the savedd data set. Thus, in one pass, you can create an outdd data set with the records that meet your specified criteria, and a savedd data set with the records that do not meet your specified criteria. DISCARD(savedd) can be used to save the records discarded by ALLDUPS, NODUPS, HIGHER(x), LOWER(y), EQUAL(v), FIRST or LAST.

Installation and Run-Time Option Enhancements

New EXPRES=p%, EXPOLD=p% and EXPMAX=p% installation options and a new HIPRMAX=p% installation and run-time option are now available. p% can be used to vary the limit DFSORT calculates for the corresponding option as a percentage of the configured expanded storage on the system at run time. If the co.figured expanded storage on a system changes, p% will cause a corresponding change in the run-time limit calculated for the corresponding option. When sharing DFSORT installation options between systems, such as in a sysplex, p% can be used to tailor the limit DFSORT calculates for the corresponding option to the system on which the application runs.

A new SPANINC installation and run-time option allows you to specify what you want DFSORT to do if it detects incomplete spanned records. This gives you control over the action (continue by eliminating incomplete spanned records and recovering valid records, or terminate), type of message (informational or error) and return code (0, 4 or 16) for incomplete spanned records.

A new OVFL0 installation and run-time option allows you to specify what you want DFSORT to do when BI, FI, PD or ZD summary fields overflow. This gives you control over the action (continue or terminate), type of message (informational or error) and return code (0, 4 or 16) for summary overflow.

A new PAD installation and run-time option allows you to specify what you want DFSORT to do when the SORTOUT LRECL is larger than the SORTIN/SORTINnn LRECL. This gives you control over the action (continue or terminate), type of message (informational or error) and return code (0, 4 or 16) for LRECL padding.

A new TRUNC installation and run-time option allows you to specify what you want DFSORT to do when the SORTOUT LRECL is smaller than the SORTIN/SORTINnn LRECL. This gives you control over the action (continue or terminate), type of message (informational or error) and return code (0, 4 or 16) for LRECL truncation.

The IBM-supplied default for ICEMAC option DSA has been changed from 16 MB to 32 MB.

The maximum value for ICEMAC option OVERRGN has been changed from 64 KB to 16128 KB.

Other Enhancements

New messages ICE178I and ICE179A provide information about reallocation of VIO work data sets.

The option-in-effect messages (ICE127I-ICE133I) are now printed for Blockset copy and merge applications.

The user exit address constant can now be passed to E32 user exits for Blockset merge applications.

Null segments in variable spanned input records are now processed by DFSORT and no longer result in termination. A null segment means that there are no more segments in the block.

OS/390 and MVS/ESA Only

DFSORT Release 14 only supports the OS/390 and MVS/ESA environments. MVS/XA and VIRTDSP processing for MVS/XA are no longer supported.

Release 13 - PTFs

New Device Support

The IBM 3590 Magnetic Tape Subsystem is supported for input, output and work data sets.

Year 2000 Features

New Y2C, Y2S, Y2Z, Y2P, Y2D and Y2B formats, in conjunction with a new Y2PAST installation and run-time option, allow you to handle two-digit year data in the following ways:

- ī Set the appropriate century window for your applications (for example, 1915-2014 or 1950-2049).
- ī Order two-digit character, zoned decimal, packed decimal, decimal or binary year data according to the century

window using Blockset SORT or MERGE (for example, order 96 representing 1996 before 00 representing 2000 in ascending sequence, or order 00 before 96 in descending sequence).

- ï Transform two-digit character, zoned decimal, packed decimal, decimal or binary year data to four-digit character year data according to the century window using OUTFIL OUTREC (for example, transform 96 to 1996 and 00 to 2000).

A new PD0 format allows you to order and transform parts of packed decimal fields (for example, month and day in date fields) using SORT, MERGE and OUTFIL.

FREE=CLOSE support for DFSPARM makes it possible to override the SORT statements generated by multiple COBOL SORT verbs in the same COBOL program.

OS/390 Registration

With OS/390 R2 and above, a check is performed to ensure that the DFSORT product is licensed for use, either as a feature of OS/390 or as a separate program product.

Performance Improvements

Performance improvements for FLR and VLR Blockset sorts include the following: ï

Dataspace sorting can now be used for variable-length record sort applications. ï

DFSORT data processing methods have been improved.

- ï Dynamic storage adjustment is a new feature that allows DFSORT to automatically use more storage than the TMAXLIM value for a Blockset sort application if DFSORT determines that doing so should improve performance. New installation options DSA=n has been added to enable you to specify the dynamic storage adjustment limit.
- ï The upper limit for the amount of main storage that can be specified and used by DFSORT has been raised from 32M to 2000M. Specifying more main storage can provide the following benefits:
 - ñ It allows DFSORT to sort very large data sets more efficiently.
 - ñ It allows more sort applications to be done entirely in main storage, eliminating the need for intermediate work space and greatly reducing the EXCP counts for those applications.
 - ñ It increases the maximum amount of data DFSORT can process in a single sort application.
- ï New installation option IOMAXBF=n has been added to enable you to specify the upper limit for the amount of storage to be used for SORTIN and SORTOUT data set buffers, which in turn limits the amount of data that can be transferred in a single I/O operation.
- ï The upper limit for the number of JCL and dynamically allocated work data sets that can be specified and used by DFSORT's Blockset technique has been raised from 32 to 100. The use of more work data sets increases the maximum amount of data that DFSORT can process in a single sort application.
- ï Changes to the DFSORT SVC provide caching selection enhancements that improve storage control caching performance, especially for SORTIN and SORTOUT devices.
- ï DFSORT can now use NOEQUALS for VLR Blockset applications if EQUALS=NO is specified at installation or NOEQUALS is specified at run-time. The use of NOEQUALS can improve performance and is recommended for applications for which the order of records that collate identically need not be preserved from input to output. To minimize migration concerns, the IBM-supplied default for the ICEMAC EQUALS option is the new value VLBLKSET, which is equivalent to EQUALS=YES for VLR Blockset applications and to EQUALS=NO for all other applications.

Floating Point for SUM

FL format can now be used with the SUM control statement for short (4-byte), long (8-byte) and extended (16-byte) floating point data.

Security Improvements

Changes to the DFSORT SVC provide security improvements that are designed to bring DFSORT up to B1 security standards.

EXCPVR Processing Removed

To enhance DFSORT's protection of system integrity, EXCPVR processing will no longer be used. EXCPVR parameter values will continue to be accepted, but will have no effect on DFSORT processing. In general, the performance improvements provided by EXCPVR processing have diminished with newer technologies and will be more than offset by the performance improvements listed above.

Release 13 - Base (May, 1995)

DFSORT's Performance Booster for The SAS System

DFSORT Release 13 provides significant CPU time improvements for SAS applications. To take advantage of this new feature, contact SAS Institute Inc. for details of the support they provide to enable this enhancement.

Dynamic Hipersorting

Dynamic Hipersorting is a new, automatic feature that eliminates the unintended system paging activity and expanded storage and paging data set space shortages that sometimes resulted from a large amount of Hipersorting activity, especially from multiple concurrent Hipersorting applications. Dynamic Hipersorting allows for more optimal DFSORT and system performance and provides installation options that allow you to customize HIPRMAX=OPTIMAL to your own criteria. With the advent of this feature, we recommend that you use HIPRMAX=OPTIMAL as your site default.

Performance

Performance enhancements for DFSORT applications that use the Blockset technique include the following:

- Improved data processing methods for fixed-length record sort applications.
- OUTFIL processing for producing multiple output data sets using a single pass over one or more input data sets.

OUTFIL Processing

OUTFIL is a new DFSORT control statement that allows you to create one or more output data sets for a sort, copy, or merge application from a single pass over one or more input data sets. You can use multiple OUTFIL statements, with each statement specifying the OUTFIL processing to be performed for one or more output data sets. OUTFIL processing begins after all other processing ends (that is, after processing for exits, options, and other control statements). OUTFIL statements support a wide variety of output data set tasks, including:

- ï Creation of multiple output data sets containing unedited or edited records from a single pass over one or more input data sets.
- ï Creation of multiple output data sets containing different ranges or subsets of records from a single pass over one or more input data sets. In addition, records that are not selected for any subset can be saved in another output data set.
- ï Conversion of variable-length record data sets to fixed-length record data sets.
- ï Sophisticated editing capabilities such as hexadecimal display and control of the way numeric fields are presented with respect to length, leading or suppressed zeros, symbols (for example, the thousands separator and decimal point), leading and trailing positive and negative signs, 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.
- ï Selection of a character or hexadecimal string for output from a lookup table, based on a character, hexadecimal, or bit string as input (that is, lookup and change).
- ï Highly detailed three-level (report, page, and section) reports containing a variety of report elements you can specify (for example, current date, current time, page number, character strings, and blank lines) or derive from the input records (for example, character fields, edited numeric input fields, record counts, and edited totals, maximums, minimums, and averages for numeric input fields).

National Language Support

ï Cultural Sort and Merge

DFSORT will allow the selection of an active locale at installation or run time and will produce sorted or merged records for output according to the collating rules defined in the active locale. This provides sorting and merging for single- or multi-byte character data based on defined collating rules which retain the cultural and local characteristics of a language.

ï Cultural Include and Omit

DFSORT will allow the selection of an active locale at installation or run time and will include or omit records for output according to the collating rules defined in the active locale. This provides inclusion or omission for single- or multi-byte character data based on defined collating rules which retain the cultural and local characteristics of a language.

ï UTFIL Reports

UTFIL allows date, time, and numeric values in reports to be formatted in many of the notations used throughout the world.

ï ICETOOL Reports

ICETOOL's DISPLAY operator allows date, time, and numeric values in reports to be formatted in many of the notations used throughout the world.

ICETOOL Enhancements

ICETOOL is now even more versatile as a result of enhancements to the existing operators. The improvements to ICETOOL include:

- ï Allowing more data to be displayed with the DISPLAY and OCCUR operators. DISPLAY now allows up to 20 fields (increased from 10) and a line length of up to 2048 characters (increased from 121). OCCUR now allows a line length of up to 2048 characters (increased from 121).

- ï More extensive formatting capabilities for numeric fields with the DISPLAY operator. Formatting items can be used to change the appearance of individual numeric fields in reports with respect to separators, decimal point, decimal places, signs, division, leading strings, floating strings and trailing strings. Thirty-three pre-defined editing masks are available for commonly used numeric editing patterns, encompassing many of the numeric notations used throughout the world. Leading and trailing strings can also be used with character fields.
- ï Display of the four-digit or two-digit year with the DISPLAY and OCCUR operators.
- ï Division of reports into sections with the DISPLAY operator, base on the values in a character or numeric break field. Statistics (total, maximum, minimum and/or average) can be displayed for each section as well as for the entire report.
- ï Automatic use of OUTFIL processing for a list of TO ddnames with the COPY and SORT operators, resulting in creation of multiple T(output) data sets from a single pass over the FROM (input) data set.
- ï Allowing OUTFIL statements to be specified in the USING data set in addition to or instead of the TO operand with the COPY and SORT operators.
- ï Allowing the active locale to be specified for the COPY, COUNT and SORT operators, in order to override the installation default for the active locale. Thus, multiple active locales can be used in the same ICETOOL job step for these operators.
- ï Allowing the last record for each unique field value to be kept with the SELECT operator.

INCLUDE/OMIT Substring Search

INCLUDE and OMIT function enhancements provide powerful substring search capability to allow inclusion or omission of records when:

- ï a specified character or hexadecimal constant is found anywhere within a specified input field (that is, a constant is a substring within a field), or
- ï a specified input value is found anywhere within a specified character or hexadecimal constant (that is, a field is a substring within a constant).

SMF Type-16 Record Enhancements

New fields, such as information pertaining to each DFSORT run about SORTIN, SORTINnn, SORTOUT and OUTFIL data sets, control statements, record counts, specified values for E15, E35, HIPRMAX, DSPSIZE, FILSZ, LOCALE and AVGRLEN, have been added to DFSORT's SMF type-16 record.

SMF=FULL, SMF=SHORT, and SMF=NO can now be specified in an OPTION statement in DFSPARM or the extended parameter list, to produce or suppress the SMF type-16 record for an individual application.

Note: The offsets of fields ICESPGN, ICEUSER, and ICEGROUP have changed in the Release 13 SMF record. If you have programs that reference those fields, recompile them using the Release 13 version of the ICESMF macro, before attempting to run them against Release 13 SMF records.

Other Enhancements

Several ICEMAC installation options have been added or changed:

- ï The IBM-supplied default for EXCPVR has been changed from ALL to NONE.
- ï The IBM-supplied default for DYNAUTO has been changed from NO to YES.
- ï SDBMSG enables you to specify whether DFSORT should use the system-determined optimum block size for DFSORT message data sets and ICETOOL message and list data sets.

- LOCALE enables you to select an active locale.
- ODMAXBF enables you to specify the maximum buffer space DFSORT can use for each OUTFIL data set.
- EXPMAX enables you to specify the maximum total amount of available storage to be used for all Hipersorting applications.
- EXPOLD enables you to specify the maximum total amount of old expanded storage to be used at any one time by all Hipersorting applications.
- EXPRES enables you to specify the minimum amount of available expanded storage to be reserved by DFSORT for use by non-Hipersorting applications.

Several run-time options have been added or changed:

- LOCALE enables you to select an active locale.
- SMF enables you to specify whether DFSORT is to produce SMF type-16 records.
- ODMAXBF enables you to specify the maximum buffer space DFSORT can use for each OUTFIL data set.
- NZDPRINT enables you to indicate that positive ZD summation results are not to be converted to printable numbers (overrides ZDPRINT).
- HILEVEL=YES on the MODS statement enables you to indicate that the E15 and E35 routines are to be treated as COBOL exits.
- DEBUG options BUFFERS=ANY and BUFFERS=BELOW will now be recognized but not used.

DFSORT will now ignore any DD statements not needed for the application (for example, a SORTIN DD statement will be ignored for a merge application).

For unsuccessful completion due to an unsupported operating system, DFSORT, ICEGENER, and ICETOOL will now pass back a return code of 24 to the operating system or invoking program.

The installation initialization exit, ICEIEXIT, enables you to specify the maximum buffer space DFSORT can use for each OUTFIL data set.

The installation termination exit, ICETEXIT, contains additional fields such as a flag to indicate that OUTFIL processing was used.

For INREC and OUTREC:

- The upper limit for columns and the end of fields has been raised from 32000 to 32752.
- 1: before the RDW field of variable-length records will be accepted and ignored.

For INCLUDE and OMIT, COND=ALL, COND=(ALL), COND=NONE, and COND=(NONE) enable you to include or omit all records.

The L2 value from the RECORD statement will be used if the L1 value is not specified when an E15 or E32 user exit passes all of the input records.

When input is a VSAM data set and output is a non-VSAM data set with RECFM not specified, DFSORT will now set the output RECFM as blocked rather than unblocked, when doing so will allow the use of the system-determined optimum block size for output.

Release 12 - PTFs

New Device Support

Four-digit device numbers are supported.

The IBM 3390-9 DASD is supported for input, output, and work data sets, although it is not recommended for work data sets for performance reasons.

The IBM RAMAC Array DASD and RAMAC Array Subsystem are supported for input, output, and work data sets.

The IBM 3990 Model 6 control unit is supported.

The IBM cached 9343 control unit models are supported.

Other Enhancements

ICEGENER, copy, and Blockset sort and merge can now be used when a tape output data set is specified with DISP=MOD or DISP=OLD, without specifying the RECFM, LRECL, or BLKSIZE in the DD statement.

Sequential striping is supported for input and output data sets.

Compression is supported for input and output data sets.

BatchPipes/MVS input and output pipes are supported.

Release 12 - Base (September, 1992)

Performance

Performance enhancements for fixed-length record sorts that use the Blockset technique include:

- Dataspace sorting, a new DFSORT capability that uses data space available with MVS/ESA on ESA systems.
- Virtual dataspace, a new DFSORT capability that enables dataspace sorting on MVS/XA systems by simulating data space.
- Improved DFSORT data processing methods.

MVS/ESA and MVS/XA Only

Beginning with Release 12, DFSORT does not run in the MVS/370 environment; only the MVS/ESA and MVS/XA environments are supported.

ICETOOL

ICETOOL, introduced in R11.1, is now even more versatile as a result of many enhancements to the DISPLAY operator, and the addition of three new operators: DEFAULTS, OCCUR, and SELECT. ICETOOL's new capabilities include:

- Creating a list data set showing the DFSORT installation defaults selected at your site.
- Creating list data sets showing character (up to 80 bytes) and numeric fields in a variety of report formats, allowing control of title, date, time, page numbers, headings, lines per page, field formats, and total, maximum, minimum and average values for the columns of numeric data.

- ï Creating list and output data sets for records with duplicate values, non-duplicate values, or values that occur n times, less than n times, or more than n times.
- ï Creating list data sets showing unique values for selected character and numeric fields and the number of times each occurs, in a variety of report formats.

Floating Sign Format

The new floating sign data format (CSF or FS) can be used with DFSORT's SORT, MERGE, INCLUDE and OMIT statements and with various ICETOOL operators to handle many data patterns that DFSORT could not previously process (for example, floating + and - signs, no sign, suppressed or specified leading zeros, and so on). This format is well suited for processing data created by many types of FORTRAN, PL/I and COBOL output formats.

For example, the following FORTRAN output data values can be processed as floating-sign data, illustrating the flexibility of this new format:

I4	SP,I4	S,I4.3	SP,I4.3
5	+5	005	+005
-247	-247	-247	-247
25	+25	025	+025
800	+800	800	+800
-4	-4	-004	-004

Panels

The DFSORT Japanese Panels now contain the enhancements made to the DFSORT English Panels in Release 11.1.

Other Enhancements

Very large values are now handled in appropriate options (up to 15 significant digits instead of 8 digits) and messages (up to 20 digits instead of 10 digits).

Several ICEMAC installation options have been added or changed: .

- ï ALTSEQ=TABLE enables you to supply a complete 256-byte alternate collating sequence table.
- ï DIAGSIM enables you to simulate the use of a SORTDIAG DD statement for DFSORT applications.
- ï DSPSIZE and VIRTDSP control the use of daspace sorting and virtual daspace, respectively.
- ï GNPAD and GNTRUNC specify the action to be taken when ICEGENER is used and the SYSUT2 LRECL is different from the SYSUT1 LRECL.

Several run-time options have been added or changed:

- ï AVGRLEN and the Un value of FILSZ and SIZE give you more control over the amount of work space DFSORT dynamically allocates.
- ï DSPSIZE controls the use of daspace sorting.
- ï nM values, for appropriate options, enable you to specify large numbers with less effort.
- ï ZDPRINT converts positive zoned-decimal sum fields into printable characters.

Message ICE071A, issued when a user exit returns an invalid return code, contains additional information to help you fix the user exit.

The type-16 SMF record contains additional fields, such as dataspace sorting statistics, Hipersorting statistics, input and output data set names and volumes, and 64-bit fields for number of records and bytes sorted.

The installation initialization exit, ICEIEXIT, enables you to control the use of dataspace sorting.

The installation termination exit, ICETEXIT, contains additional fields, such as dataspace sorting statistics.

Specifying a MAINSIZE (or SIZE) value equal to your MAXLIM value is no longer equivalent to specifying MAINSIZE=MAX (or SIZE=MAX) and results in RESALL and RESINV values of zero.

Release 11.1 - PTFs

New Device Support

Support for the following devices was made available:

- ï IBM 3990 ESCON channels and non-synchronous control units
- ï IBM 9340 Direct Access Storage Subsystem, which includes the 9341 Storage Control Module, the 9343 Storage Controller and the 9345 Direct Access Storage Device
- ï IBM 3490 Magnetic Tape Subsystem Enhanced Capability (3490E) Models A10, A20, B20, B40, C10, C11, C22, D41, and D42
- ï IBM 3490 Magnetic Tape Subsystem ESCON Adapter
- ï IBM Enhanced Capacity Cartridge System Tape for 3490E.

Other Enhancements

INCLUDE and OMIT statement enhancements provide powerful bit-level logic capability using two different methods:

- ï Bit operator with hexadecimal or bit mask
- ï Bit comparison with bit constant

An enhancement to DFSORT, in conjunction with DFP APAR OY31441, now enables IDCAMS BLDINDEX to automatically use DFSORT to improve the performance of most BLDINDEX jobs that currently use the BLDINDEX external sort. You do not need to change any BLDINDEX jobs, and you can now sort indexes in an all-DFSMS environment.

Short-record processing with the VLSHRT option is now supported for merge applications that use the Blockset technique. Removing the NOBLKSET option from merge applications with short records can result in improved performance.

The file size for multivolume and concatenated input data sets can now be determined automatically, making it easier to use dynamically allocated work data sets.

The requirement for the data set with the largest blocksize to be specified first, when concatenated input includes tape data sets for a sort application, has been removed.

The blocksize for message data sets is no longer limited to 121; any valid blocksize can be used.

The VSAM variable-length RRDS (VRRDS) attribute is supported.

DFSORT, in conjunction with DFP APAR OY50285, sets the reblockable indicator in the output data set label when DFSORT or allocation sets the system-determined optimum blocksize.

Several ICEMAC installation options have been added:

- ï CFW controls the use of cache fast write for work data sets. A run-time option is also available for this purpose.
- ï DYNPC controls dynamically allocated work space when the input file size is unknown to DFSORT.
- ï IDRCPC controls dynamically allocated work space for input data sets using the IDRC feature of 3480 and 3490 devices.
- ï VSAMBSP enables you to improve DFSORT's performance for sort applications through the use of additional VSAM buffers.

Release 11.1 - Base (February, 1991)

ICETOOL

ICETOOL is a versatile new DFSORT utility that allows you to perform multiple operations on one or more data sets in a single job step. The nine ICETOOL operators, each of which can be used one or more times in a single run, allow DFSORT users to perform a variety of functions, such as:

- ï Displaying statistical information for selected numeric fields such as minimum, maximum, average, total, count of values within a range and count of unique values. This makes it easy to extract frequently used analytical data.
- ï Displaying selected numeric fields in readable format along with character fields. This makes it easy to view numeric fields which are normally in unprintable format.
- ï Identifying and displaying invalid decimal values and their locations in the data set. This makes it easy to avoid using invalid fields for other operations.
- ï Creating multiple copies of sorted, edied, or unedited data sets. This makes it easy to create several identical data set copies.
- ï Creating output data sets containing different subsets of input data sets or field arrangements of input data sets. This makes it easy to view data in many different ways.
- ï Allowing operations to be performed or suppressed based on the success or failure of previous operations. This makes it easy to group operations according to the action to be taken after an error.

ICETOOL can be called directly or from a program. Operations can be supplied in a data set or in a parameter list passed by a calling program. For each operator supplied in the parameter list, ICETOOL returns information in the parameter list pertaining to that operation, thus allowing the calling program to use the information derived by ICETOOL.

Panels

Enhancements to the DFSORT English panels increase user productivity:

- ı The full range of DFSORT control statement features is now available through panels.
- ı New panels allow you to allocate new output data sets for your applications.
- ı You can now generate control statements in expanded form, making them easier to read and to edit.
- ı LIBDEF support is now available if you wish to define DFSORT as an application-level set of libraries to be searched before the allocated ISPF libraries.

Other Enhancements

For sort and copy applications, short record processing with the VLSHRT option is now supported for compare fields specified in an INCLUDE or OMIT statement.

Several new installation options are supported by ICEMAC:

- ı FSZEST=YES/NO allows you to specify whether or not DFSORT treats exact filesize values as estimated filesize values at run-time.
- ı TSO specifies the set of ICEMAC defaults used when DFSORT is invoked directly by foreground TSO users.
- ı TSOINV specifies the set of ICEMAC defaults used when DFSORT is invoked from programs by foreground TSO users.

The following alternate forms of existing operands are allowed:

```
SORT      FIELDS=(COPY)
MERGE FIELDS=(COPY) SUM
FIELDS=(NONE)
```

The limit for the number of dynamically allocated work data sets (DYNALOC/DYNALLOC) is raised from 16 to 32.

The extended parameter list now supports a call identifier field which can be used by invoking programs to identify the messages associated with individual calls to DFSORT.

The 24-bit parameter list no longer requires a blank at the end of control statements and now supports the following additional entries for control statements:

```
INCLUDE/OMIT
OUTREC INREC
OUTFIL - accepted but not processed.
```

Conflicting format values for SORT, MERGE, SUM, INCLUDE, and OMIT no longer result in termination. If valid format values are specified in both FORMAT and FIELDS or COND, DFSORT issues an informational message, uses the format values from FIELDS/COND, and ignores the format value from FORMAT.

Performance

Improved performance through internal enhancements to the Blockset technique (fixed-length record sorting only).

Release 11 - PTFs

New Device Support

Support for IBM 3390 Direct Access Storage Device has been added.

DFSORT supports all models of the ES/9000 processor. DFSORT applications based on the System/370 architecture will run effectively on ES/9000 and MVS/ESA SP Version 4 with no modification.

Other Enhancements

DYNAUTO=IGNWKDD is a new ICEMAC option that allows you to use dynamic allocation of work data sets at run-time even if JCL work data sets are present. OPTION USEWKDD in DFSPARM can be used to override DYNAUTO=IGNWKDD for individual jobs.

DFSORT now deallocates dynamically allocated work data sets at the end of an application, removing the requirement that the largest of multiple sorts be done first.

The data set with the largest block size can be specified anywhere in the SORTIN concatenation for copy and DASD work data set sort applications.

Support for Partitioned Data Set Extended (PDSE) has been added.

ICEMAC option SDB=YES/NO allows you to specify whether DFSORT should use the system-determined optimum block size for DASD and tape SORTOUT when appropriate.

Release 11 - Base (November, 1988)

Enhancements

Performance enhancements to the Blockset technique include:

- ï Hipersorting, a new DFSORT capability that uses hiperspace available with MVS/ESA on Enterprise Systems Architecture/370.
- ï New internal algorithms and more extensive use of System/370-XA sorting instructions for MVS/XA and MVS/ESA (fixed-length record record sorting only).
- ï Use of EXCPVR when sorting and copying data sets for MVS/XA and MVS/ESA.
- ï More efficient allocation of dynamic work data sets.

New installation and run-time parameters are available to modify the use of:

- ï Hipersorting
- ï EXCPVR
- ï Control interval access for VSAM data sets
- ï DFSORT's ESTAE recovery routine.

PARM options and DFSORT control statements can now be supplied to DFSORT from a single source data set for JCL and program-invoked DFSORT jobs.

The INREC and OUTREC statements have been enhanced:

- ı Character and hexadecimal string constants are now supported as separators.
- ı Column alignment is now supported for separators and input fields.
- ı The upper limit for the separation repetition factor is raised from 256 to 4095.

The Conventional technique modules can be installed in a system library to eliminate the need for the SORTLIB DD statement.

Duplicate DD statements are ignored after the first-specified DD statement.

SORTKW0 through SORTWK9 DD statements are now accepted and treated as duplicates of SORTWK00 through SORTWK09 DD statements. SORTIN0 through SORTIN9 DD statements are now accepted and treated as duplicates of SORTIN00 through SORTIN09 DD statements.

DFSORT no longer requires that the SORTINnn data set with the largest block size and record length be specified first for Blockset merge jobs.

The maximum length for variable blocked spanned records is increased to 32767.

Multivolume SORTWKnn data sets are now allowed, but only the first volume of each data set is used.

DFSORT can now handle more complex SUM applications.

The MINLIM default is raised from 200K to 440K.

Release 10 - Base (January, 1988)

Enhancements

Better performance through internal enhancements to the Blockset technique, which reduce the required SORTWK space, CPU time, and EXCPs.

DFSORT Panels, an interactive menu-driven facility supported under ISPF and ISMF, lets you use DFSORT on-line with minimal experience.

New ICEMAC options:

- ı With DYNAUTO, you can specify automatic dynamic allocation of intermediate data sets.
- ı With ABCODE, you can customize critical error ABEND codes using either the critical message numbers as before, or your own choice of a number between 1 and 99.
- ı With EXITCK, you can prevent termination when E15/E35 user routines use certain "invalid" return codes.
- ı With NOMSGDD, you can prevent termination when a needed message data set is not present.

The MODS statement allows you to specify that user exit routines are in the JOBLIB/STEPLIB or link library data sets.

A new subparameter of the DYNALLOC execution option lets you suppress dynamic allocation of intermediate data sets.

New parameters added to the EXEC statement let you specify options formerly available only on ICEMAC or control statements. Aliases (for example, MSGDD for MSGDDN) are also available.

For increased flexibility, DFSORT now accepts:

- ï Semicolons interchangeably with commas in the EXEC PARM string and within control statements.
- ï Blank statements.
- ï FB, VB and DB for RECORD TYPE.
- ï Duplicate message data set DD statements (the first is used).
- ï SORTWKxx DD statements for a merge applications (ignored).
- ï Empty VSAM input data sets for sort and copy applications.
- ï VSAM input and non-VSAM output without the RECORD TYPE parameter.
- ï Shorter output LRECL for VLR data sets.
- ï The STOPAFT parameter on the SORT and MERGE statement.
- ï LS, TS, OL, and OT as format types on the SORT, MERGE, INCLUDE, and OMIT statements. Automatic

estimation of VSAM input data set size for intermediate storage optimization.

DFSORT uses the job stepname rather than the procedure stepname in messages and SMF field ICESTPNM.

Release 9 - Base (April, 1987)

Enhancements

Improved performance:

- ï For sort, merge and copy applications from internal enhancements to the Blockset technique.
- ï For fixed-length, nonspanned, VSAM data sets from more efficient input processing in the Blockset technique.

Installation-written exit routines (ICEIEXIT and ICETEXIT) can be used to exercise control over individual jobs and to collect job statistics.

Use of the more efficient DFSORT copy function for qualifying IEBGENER jobs (ICEGENER).

Improved capability to run two different releases of DFSORT at the same time.

Additional statistical information in the SMF record.

Capability to make zoned-decimal positive-sum fields printable.

User written Extended Function Support (EFS) programs can be used to:

- ï Examine, alter, and/or ignore DFSORT and non-DFSORT control statements or EXEC PARM options.
- ï Support sorting or merging of user-defined data types with user-defined collating sequences.
- ï Provide the logic to include or omit records based on user-defined data types.
- ï Provide user supplied messages to DFSORT for printing to the message data set.

MVS/ESA and MVS/XA only

DFSORT Release 9 only supports the MVS/ESA and MVS/XA environments. OS/VS1 is no longer supported.

Release 8 - Base (March, 1986)

Enhancements

An enhancement to the Blockset technique can improve performance when sorting variable-length records. On MVS/XA systems, utilization of extended addressing capability is available for VLR record sorts.

More efficient use of processor cache memory can improve performance when sorting fixed-length records.

The new copy function copies a data set without performing any sorting or merging operation. Copy can be used with most of the same control statements, exits and options available when sorting or merging.

New installation and run-time options can now be used to continue sorting when a record too short to contain all specified control fields is encountered.

The Blockset technique now supports VSAM data sets for input and output.

Release 7.1 - Base (June, 1985)

Enhancements

Preserve the original sequence of identically collating records when doing a Blockset merge if the EQUALS option is used.

Specify, using the EQUALS option, that the first record will be retained when summarizing identically collating records for a Blockset sort or merge.

Use the Blockset technique when merging spanned variable-length records.

Release 7 - Base (January, 1985)

New Device Support

The 3480 Magnetic Tape Subsystem is supported for intermediate storage.

Other Enhancements

For MVS/XA, reduce the processor time for sorting in large storage by using IBM System/370 Extended Architecture Sorting Assists.

Virtual storage constraint relief for MVS/XA:

- ı If installing DFSORT resident, most of the fixed-length record Blockset technique modules will be placed above 16-megabyte virtual in the extended link pack area. This provides more space in the link pack area for resident programs that cannot reside above 16-megabyte virtual.
- ı When sorting fixed-length records using the Blockset technique, DFSORT can place selected buffers above 16-megabyte virtual, leaving more space below 16-megabyte virtual for user applications.
- ı Performance is enhanced by using more and larger buffers above 16-megabyte virtual.
- ı The ability to specify:
 - ı Upper limit of the amount of main storage above and below 16-megabyte virtual available to DFSORT (TMAXLIM).
 - ı Number of bytes reserved above 16-megabyte for system use (ARESALL).
 - ı Number of bytes reserved above 16-megabyte virtual for a program that invokes DFSORT (ARESINV).

COBOL-related enhancements:

- ı DFSORT can be invoked from VS COBOL II programs.
- ı The VS COBOL II FASTSORT compile time option can improve performance.
- ı E15 and E35 user exit routines can be written in COBOL.
- ı An alternate message data set can be specified when invoking DFSORT with JCL. This is especially useful when exits are written in COBOL.

The upper limit of 48K for the RESALL option has been removed.

The maximum number of records to be accepted for sorting can now be specified.

Release 6 - Base (January, 1984)

Enhancements

Enhancements to the Blockset technique can improve performance for sort and merge applications.

DFSORT can now interface with programs and exits that use the new addressing mode and residence mode capabilities of MVS/XA.

The following control statements have been added:

- ı INCLUDE and OMIT to allow inclusion or omission of records, based on logical expressions.
- ı INREC and OUTREC to allow modification of input or output record layouts.
- ı SUM to allow summarization of records with equal control fields.

An extended parameter list can be used to invoke DFSORT from another program.

A 31-bit user address constant can be passed to exits E15 and E35.

Options can be specified from several different sources to allow more flexibility in overriding installation options.

Additional parameters can be specified as EXEC PARMs and in the OPTION statement.

The number of input files can be specified in the MERGE statement when input is supplied by exit E32.

"Options-in-effect" information is printed in the message data set.

The BSAM option can be used for Blockset.

Special diagnostic messages can be obtained by specifying a SORTDIAG DD statement.

Comment statements can be specified.

Control statement operands can be continued in any column from 2 through 71.

A maximum of 32 SORTWKnn or 16 dynamically allocated work data sets can be used.

The minimum amount of storage needed to run DFSORT has been raised to 88K.