

IBM DB2 Sort for z/OS User's Guide

Version 2 Release 1



IBM DB2 Sort for z/OS User's Guide

Version 2 Release 1

Note:

Before using this information and the product it supports, read the "Notices" topic at the end of this information.

First Edition (October 2014)

This edition applies to Version 2 Release 1 of IBM DB2 Sort for z/OS (product number 5655-AA9) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2010, 2014; Copyright Syncsort Incorporated 2010, 2014.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this information	v	Specifying the metadata library for the product to customize	36
Chapter 1. DB2 Sort overview	1	Discovering DB2 Sort information automatically	38
What does DB2 Sort do?	2	Creating and associating DB2 entries	40
Memory usage optimization	2	Defining parameters	41
zIIP technology exploitation	2	Generating customization jobs	45
Optimization of sort work space usage	2	Submitting customization jobs	45
How DB2 Sort works with DB2 utilities and tools	4	Browsing parameters	47
DB2 Sort scenarios	6	Copying DB2 entries	47
DB2 Sort components	6	Removing DB2 entries	49
Tools Customizer overview	6	Deleting DB2 entries	49
Service updates and support information	7	Displaying customization jobs	50
Product documentation and updates	7	Maintaining customization jobs	50
Accessibility features	8	Using Tools Customizer in a multiple-LPAR environment	51
Summary of changes	8	Accessing DB2 Sort libraries	52
Chapter 2. Preparing to customize DB2 Sort	11	Chapter 4. Verifying installation and customization.	53
Worksheets: Gathering required data set names and parameter values	12	Installation Verification Program: Overview.	53
Default installation options	15	Using Installation Verification Program	54
Tools Customizer terminology and data sets	18	Generating an IVP report.	54
Tools Customizer terminology	18	Running with DB2 Sort	55
Data sets that Tools Customizer uses during customization	20	Chapter 5. Troubleshooting DB2 Sort	57
Runtime considerations	23	DB2 Sort messages	57
DB2 Sort messages	23	DB2 Sort messages	57
Gathering diagnostic information	23	Gathering diagnostic information	70
Chapter 3. Activating and customizing DB2 Sort	27	Chapter 6. Tools Customizer reference	73
Activating DB2 Sort	27	Tools Customizer terminology and data sets	73
Starting and preparing Tools Customizer for use	28	Tools Customizer terminology	73
Starting Tools Customizer	28	Data sets that Tools Customizer uses during customization	75
Modifying Tools Customizer user settings	29	Notices	79
Hiding and displaying panel text	32	Index	83
Customizing DB2 Sort	33		
Roadmap: Customizing DB2 Sort for the first time	33		
Roadmap: Customizing a new version of DB2 Sort from a previous customization	34		
Roadmap: Recustomizing DB2 Sort	35		

About this information

IBM DB2[®] Sort for z/OS (also referred to as DB2 Sort) is a DB2 Tool that provides high-speed utility sort processing for data that is stored in DB2 for z/OS databases.

These topics are designed to help database administrators, system programmers, application programmers, and system operators perform these tasks:

- Activate and customize DB2 Sort
- Diagnose and resolve DB2 Sort issues

To use these topics, you should have a working knowledge of:

- The z/OS operating system
- ISPF

Always check the DB2 Tools Product Documentation page for the most current version of this information:

<http://www.ibm.com/software/data/db2imstools/db2tools-library.html>

Chapter 1. DB2 Sort overview

IBM® DB2 Sort for z/OS® (also referred to as DB2 Sort) is a DB2 Tool that provides high-speed utility sort processing for data that is stored in DB2 for z/OS databases.

DB2 Sort improves the sort performance of many of the DB2 utilities in the IBM DB2 Utilities Suite and of several other DB2 management tools. The way DB2 Sort improves sort performance is through enhanced sort technology and by communicating with DB2 utilities and tools and then adjusting system resources to ensure optimal sorting. This approach to sorting can result in significantly reduced sort CPU time.

The following DB2 utilities and DB2 Tools can benefit from improved sorting performance by using DB2 Sort:

DB2 utilities:

- CHECK DATA
- CHECK INDEX
- CHECK LOB
- LOAD
- REBUILD INDEX
- REORG TABLESPACE
- RUNSTATS

Note: DB2 Sort supports all of the utilities in the DB2 Utilities Suite for z/OS V10.1 and V11.1.

DB2 Tools:

- DB2 Utilities Enhancement Tool for z/OS V2.2 and subsequent releases
- DB2 Log Analysis Tool for z/OS V3.3 and subsequent releases
- DB2 Recovery Expert for z/OS V3.1 and subsequent releases
- DB2 Change Accumulation Tool for z/OS V3.1 and subsequent releases
- DB2 High Performance Unload for z/OS V4.2 and subsequent releases

Topics:

- “What does DB2 Sort do?” on page 2
- “Memory usage optimization” on page 2
- “zIIP technology exploitation” on page 2
- “Optimization of sort work space usage” on page 2
- “How DB2 Sort works with DB2 utilities and tools” on page 4
- “DB2 Sort scenarios” on page 6
- “Product documentation and updates” on page 7
- “Accessibility features” on page 8

What does DB2 Sort do?

DB2 Sort can improve the performance of sort processing, especially in environments with large volumes of data, large table spaces or large indexes. These environments require a more sophisticated sorting approach than the approach that is used by tools that are used for general sorting purposes.

Sort processing in these environments often involves intensive CPU usage and heavy I/O activity, which in turn requires large amounts of memory and disk space. DB2 Sort meets these demands by optimizing the use of memory, exploiting zIIP technology, and optimizing the management of sort work disk space.

Memory usage optimization

To optimize memory usage, DB2 Sort determines the availability of system resources and chooses the appropriate amount of storage. It employs memory objects and data space to avoid DASD access when possible.

DB2 Sort coordinates the use of central storage among simultaneous sorts within an address space, optimizing intra-region parallelism. DB2 Sort also monitors the use of central storage by all DB2 Sort applications across an LPAR. By monitoring the resource usage of this inter-region parallelism, DB2 Sort can allocate optimal amounts of storage while accommodating the needs of other work that is being performed on the system.

Tip:

Because DB2 Sort manages memory allocation dynamically, setting REGION=0M is recommended for best performance.

zIIP technology exploitation

Almost every DB2 Sort instance invoked by DB2 utilities and tools offloads a portion of work to a zIIP engine, if one is available. This offload reduces the billable CPU time that is associated with sorting.

Optimization of sort work space usage

Sort work data sets are intermediate disk data sets that are used when the input file is too large to fit into memory. Using DB2 Sort can help you optimize sort work space usage.

The following sections contain information about how DB2 Sort works with the DB2 for z/OS subsystem parameter **UTSORTAL** and the DB2 utility option SORTNUM.

Note: For the best resource utilization, set the **UTSORTAL** parameter to YES.

For information about setting the **UTSORTAL** parameter, see one of the following documents:

- DB2 11 for z/OS Installation and Migration Guide (GC19-4056)
- DB2 10 for z/OS Installation and Migration Guide (GC19-2974)

To understand how DB2 Sort can optimize sort work space usage, consider the following questions.

- “What is sort work space?” on page 3
- “What does DB2 Sort improve?” on page 3

- “How does DB2 Sort reduce the amount of sort work space?”
- “How does DB2 Sort avoid Sort Capacity Exceeded errors?” on page 4
- “How does DB2 Sort improve sort work I/O performance?” on page 4
- “What is the dynamic allocation retry feature?” on page 4

What is sort work space?

A sort with the least amount of I/O activity is one in which all of the input records can be held in memory simultaneously. In this case, the records can be put into the proper sequence with only one pass of the input data and can be passed directly to the output exit routine for the utility.

When the input file is too large to fit into memory, intermediate disk data sets must be used. These data sets are called sort work data sets or sort work space. Records have to be written to the sort work data sets and subsequently read back in, sometimes multiple times, until they can be passed in the proper sequence to the output exit. This process is complex and consumes time and disk space.

What does DB2 Sort improve?

The management of sort work data sets by DB2 Sort can provide significant benefits over the sort work processing that is currently available with the DB2 utilities and tools. Improvements can be made in cases in which the user controls the number of sort work data sets that are allocated, and also in cases where the DB2 utilities and tools themselves allocate the sort work data sets.

The performance of utility and tool applications that have been set up to allocate sort work data sets dynamically, either through the use of the SORTNUM parameter to specify a number of sort work data sets or through the use of the UTSORTAL parameter, can be improved in two areas:

- A potential reduction in the amount of space allocated
- More resilience in avoiding Sort Capacity Exceeded failures when the file size estimate is inaccurate

How does DB2 Sort reduce the amount of sort work space?

The DB2 SORTNUM parameter specifies the number of sort work data sets to be allocated dynamically at the start of a sort to be done by a utility. Without DB2 Sort, this space is acquired immediately, even though it might not be filled until later in processing. This parameter also limits the number of sort work data sets that can be allocated dynamically.

With DB2 Sort, when the number of sort work data sets is fewer than 32, DB2 Sort allows the maximum number of sort work data sets to grow to 32. The initial allocation starts with the smallest number of sort work allocations that is required to achieve an appropriate overall data transfer rate when data is read or written to sort work space. This space is increased incrementally as the data to be sorted is provided by the utility calling the sort. This incremental allocation consumes less space overall when measured over the duration of the sort processing. Although the total amount of disk space used can be the same in either technique, the disk space will not have been held for as long because the space was added incrementally. This approach allows this disk space to be available in the temporary pool for other jobs for a longer period of time. With the mix of work in a batch environment, the temporary space pool can maintain a higher free space average, which helps to support other work that is running on the system.

This incremental allocation technique also addresses problems that can arise when the file size that is estimated by a DB2 utility or tool is inaccurate. If the file size is overestimated, sort work space can be acquired but never used. This space is unavailable for other applications on the system. With the incremental allocation strategy provided by DB2 Sort, this disk space is never acquired, which conserves the resource and the CPU time that would be spent in allocation.

How does DB2 Sort avoid Sort Capacity Exceeded errors?

A more serious condition arises when the file size is underestimated. In this case, the existing technique does not have the ability to allocate more sort work data sets. This condition can result in a Sort Capacity Exceeded error and the need to rerun the job with an amended SORTNUM parameter. DB2 Sort has the ability to continue to allocate additional sort work data sets. This capability makes it possible to avoid the Sort Capacity Exceeded error and allow the utility or tool to complete.

In another scenario, the utility or tool can be run with the UTSORTAL parameter set to YES instead of using the SORTNUM parameter. In this case, the utility or tool itself can dynamically pre-allocate the sort work data sets if none are provided. Although, in this case DB2 Sort cannot help with incremental allocation because the disk space is being provided, it can help when the file size has been underestimated. This situation might result in a Sort Capacity Exceeded error. DB2 Sort continues to incrementally allocate additional sort works as needed to avoid this error condition.

How does DB2 Sort improve sort work I/O performance?

DB2 Sort can add more sort work data sets if the currently available set of sort work data sets is not meeting the performance objectives for data transfer rate. The addition of sort work data sets at this point can improve the bandwidth of the write operations, which helps DB2 Sort to meet its elapsed time performance objectives. DB2 Sort automatically releases any extra sort work space at the end of the input portion of the sort process. This capability allows DB2 Sort to hold only the sort work space that is actually needed during the output portion of the sort process.

What is the dynamic allocation retry feature?

DB2 Sort also provides resiliency during its dynamic allocation operation with its ability to retry a dynamic allocation request after the request has failed due to a space shortage in the DASD temporary work pool. This retry facility can retry the dynamic allocation request a prescribed number of times with a specified amount of time between attempts. This capability helps to ensure that sorts have the best opportunity for a successful completion without terminating with a Sort Capacity Exceeded condition.

How DB2 Sort works with DB2 utilities and tools

DB2 Sort communicates with DB2 utilities and tools and then dynamically allocates system resources to optimize each sort request. When multiple sorts are run in parallel, DB2 Sort chooses the most effective use of resources for each sort, which can potentially increase the number of sorts that can be run simultaneously. During the execution of these sorts, DB2 Sort also monitors and adjusts the allocation of system resources to optimize CPU processing, I/O performance, and memory usage.

Note: To activate DB2 Sort, the DB2 subsystem parameter **DB2SORT** must be set to ENABLE. For instructions, see the following IBM support document: <http://www-01.ibm.com/support/docview.wss?uid=swg1PM12819>.

The following figure shows the series of steps that are performed by DB2 utilities and DB2 Sort when parallel sorts are invoked.

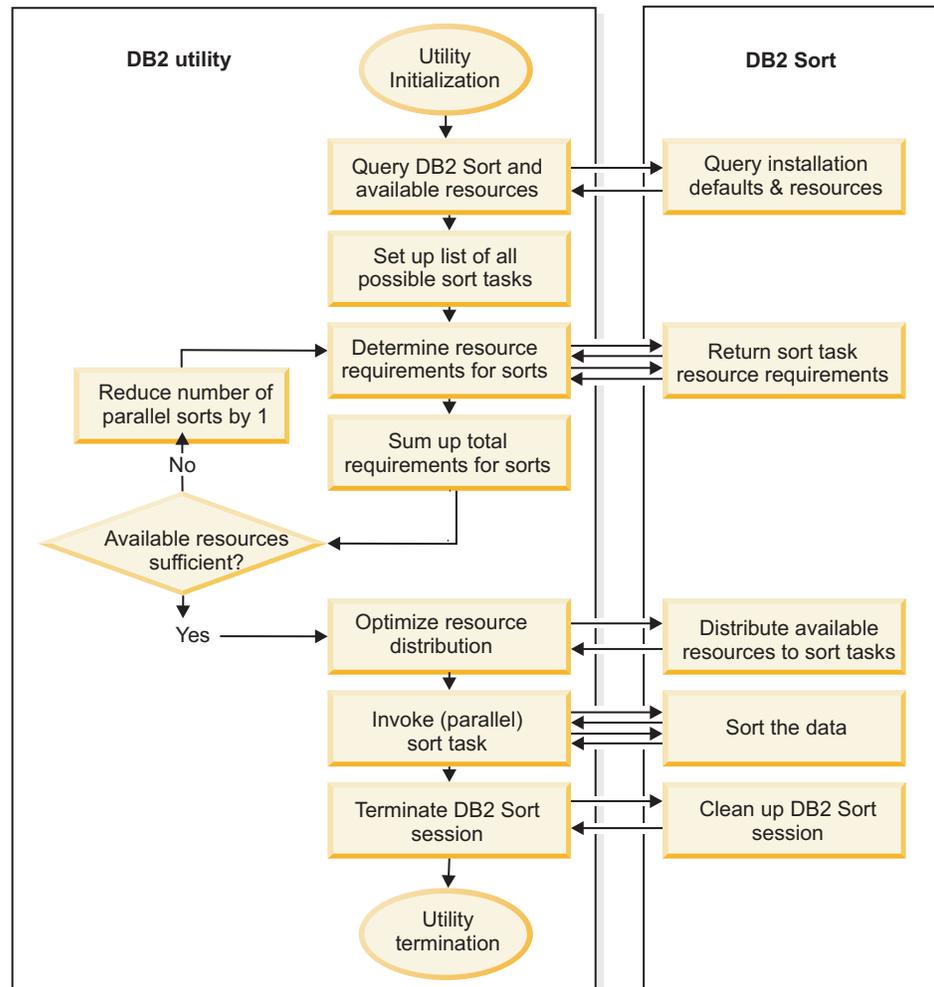


Figure 1. Process flow for DB2 Utilities and DB2 Sort

How parallel sorts are handled

When DB2 utilities require sort processing, multiple sorts from a single utility are often run in parallel to reduce the elapsed time of the application. This method of sorting is referred to as *intra-regional parallelism*. DB2 Sort communicates with the DB2 utility to choose the optimal number of sorts to run in parallel. DB2 Sort assesses the availability of system resources and communicates with the DB2 utility to optimize the allocation of resources for each sort based on the data characteristics. By choosing the most effective use of resources for each sort, DB2 Sort can increase the number of sorts that can be run simultaneously. This higher degree of parallelism can result in lower elapsed time for the application. DB2 Sort also synchronizes the processing of simultaneous sorts, reducing the wait time associated with the transfer of records between the utilities and DB2 Sort. This synchronization lowers the CPU time and elapsed time of the sort processing. The LOAD, CHECK INDEX, REBUILD INDEX and REORG TABLESPACE utilities can all benefit from the way that DB2 Sort handles parallel sorts.

DB2 Sort scenarios

You can use DB2 Sort to address many of your common database sorting problems, such as high CPU usage.

Scenario: DB2 Sort reduces high CPU usage

DB2 Sort can help reduce utility and tool sort CPU consumption overall. In addition, DB2 Sort can offload a proportion of CPU processing to zIIP processors, thereby further reducing load on general purpose processors. For example, if batch processing with many queries and updates is running simultaneously with object maintenance using REORG or RUNSTATS, DB2 Sort can enable more concurrent SQL load, which in turn can reduce the CPU usage. The result of reduced CPU usage could be either reduced internal costs or more efficient workload, or both.

DB2 Sort components

Tools Customizer is a component that is installed when DB2 Sort gets installed. Tools Customizer gives you the option to change the default installation settings.

Tools Customizer overview

IBM Tools Customizer for z/OS (also referred to as Tools Customizer) standardizes many of the customization processes that are required to customize IBM Tools that run on z/OS. Tools Customizer is a component of IBM Tools Base for z/OS.

Tools Customizer provides a consistent ISPF interface to ensure that the customization process is the same for all IBM Tools products and solution pack components. It also provides the ability to "discover" parameter values from products or solution pack components that you previously customized manually or by using Tools Customizer.

Features and benefits

Tools Customizer provides the following features:

- A single, consistent ISPF interface ensures that the customization process is the same for all IBM Tools products and solution pack components.
- A Discover EXEC discovers values for common product and DB2 parameters from a product or solution pack component that you previously customized manually or by using Tools Customizer. Each IBM Tools product and solution pack component has a unique Discover EXEC. The discovered parameters are stored in the data store. If the product or solution pack component that you want to customize exists in the Tools Customizer data store, Tools Customizer issues a warning before it overwrites existing values. Use the Discover EXEC by issuing the DISCOVER command on the Customizer Workplace panel.
- The data store retains discovered and manually specified parameter values. Because the parameter information is persistently stored, you have to manually specify or discover parameter values only once. Tools Customizer uses these parameter values where they are applicable.
- A metadata repository contains the members that define the following customization attributes for products and solution pack components:
 - Parameters, tasks, and steps for the product or solution pack component to be customized. Some product or solution pack parameters, tasks, and steps are required.

- DB2 parameters for the DB2 subsystem, DB2 group attach name, or DB2 data sharing member on which you will customize the product or solution pack component. All of the DB2 parameters are required.
- Default values are provided for product parameters and solution pack component parameters and DB2 parameters. The default values show examples of how to complete fields.

Service updates and support information

Service updates and support information for this product, including software fix packs, PTFs, frequently asked questions (FAQs), technical notes, troubleshooting information, and downloads, are available from the web.

To find service updates and support information, see the following website:

http://www.ibm.com/support/entry/portal/Overview/Software/Information_Management/DB2_Tools_for_z~OS

Product documentation and updates

DB2 Tools information is available at multiple places on the web. You can receive updates to DB2 Tools information automatically by registering with the IBM My Notifications service.

Information on the web

The DB2 Tools Product Documentation web page provides current product documentation that you can view, print, and download. To locate publications with the most up-to-date information, refer to the following web page:

<http://www.ibm.com/software/data/db2imstools/db2tools-library.html>

You can also access documentation for many DB2 Tools from IBM Knowledge Center:

<http://www.ibm.com/support/knowledgecenter>

Search for a specific DB2 Tool product or browse the **Information Management > DB2 for z/OS family**.

IBM Redbooks® publications that cover DB2 Tools are available from the following web page:

<http://www.redbooks.ibm.com>

The Data Management Tools Solutions website shows how IBM solutions can help IT organizations maximize their investment in DB2 databases while staying ahead of today's top data management challenges:

<http://www.ibm.com/software/data/db2imstools/solutions/index.html>

Receiving documentation updates automatically

To automatically receive emails that notify you when new technote documents are released, when existing product documentation is updated, and when new product

documentation is available, you can register with the IBM My Notifications service. You can customize the service so that you receive information about only those IBM products that you specify.

To register with the My Notifications service:

1. Go to <http://www.ibm.com/support/mysupport>
2. Enter your IBM ID and password, or create one by clicking **register now**.
3. When the My Notifications page is displayed, click **Subscribe** to select those products that you want to receive information updates about. The DB2 Tools option is located under **Software > Information Management**.
4. Click **Continue** to specify the types of updates that you want to receive.
5. Click **Submit** to save your profile.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other IBM product documentation, use one of the following options:

- Use the online reader comment form, which is located at <http://www.ibm.com/software/data/rcf/>.
- Send your comments by email to comments@us.ibm.com. Include the name of the book, the part number of the book, the version of the product that you are using, and, if applicable, the specific location of the text you are commenting on, for example, a page number or table number.

Accessibility features

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use a software product successfully.

The major accessibility features in this product enable users to perform the following activities:

- Use assistive technologies such as screen readers and screen magnifier software. Consult the assistive technology documentation for specific information when using it to access z/OS interfaces.
- Customize display attributes such as color, contrast, and font size.
- Operate specific or equivalent features by using only the keyboard. Refer to the following publications for information about accessing ISPF interfaces:
 - *z/OS ISPF User's Guide, Volume 1*
 - *z/OS TSO/E Primer*
 - *z/OS TSO/E User's Guide*

These guides describe how to use the ISPF interface, including the use of keyboard shortcuts or function keys (PF keys), include the default settings for the PF keys, and explain how to modify their functions.

Summary of changes

This topic summarizes the technical changes for this edition.

New and changed information is indicated by a vertical bar (|) to the left of a change. Editorial changes that have no technical significance are not noted.

Version 2.1, October, 2014, SC19-4387-00

Information has been added or revised for new enhancements and maintenance. New and changed information is marked by a vertical bar | to the left of the change. The major additions and changes include:

- Interface enhancement for DB2 Tools and Utilities that invoke DB2 Sort:
 - Reduces CPU time and elapsed time
 - Offloads more work to zIIP
- High Performance FICON channel programs for I/O to SORTIN and SORTOUT:
 - Reduce elapsed time

Chapter 2. Preparing to customize DB2 Sort

Before you start to customize DB2 Sort, determine all of the customization values that you need to specify during the customization process, and familiarize yourself with all of the customization tasks.

The following checklist lists and describes each significant customization step. Use this checklist to guide you through the entire customization process.

Tip: Print the following checklist and the data set names and parameter values worksheets. Use the worksheets to record your values, and refer to them during the customization process.

Task	Link to detailed instructions	Status
Tools Customizer basics		
Prior to beginning the customization process, familiarize yourself with Tools Customizer terminology and data sets, and other basic information about Tools Customizer.	"Tools Customizer terminology and data sets" on page 18	
SMP/E installation		
Verify that DB2 Sort has been installed correctly. DB2 Sort is installed by using standard SMP/E processing.	See <i>IBM DB2 Sort for z/OS Program Directory</i> , GI10-9004, for installation instructions.	
Verify that Tools Customizer for z/OS has been installed correctly. Tools Customizer for z/OS is installed by using standard SMP/E processing.	See <i>IBM Tools Base for z/OS Program Directory</i> , GI10-8819, for installation instructions.	
Gather data set names and parameter values		
During the customization process, you must specify names for the following types of data sets: <ul style="list-style-type: none"> • Tools Customizer • DB2 Sort 	"Worksheets: Gathering required data set names and parameter values" on page 12	
During the customization process, you must specify parameter values for DB2 Sort, and in some cases, for DB2.	"Default installation options" on page 15	
Activate DB2 Sort		
Apply appropriate APARs and set appropriate DB2 parameters.	"Activating DB2 Sort" on page 27	
Customize DB2 Sort		
Start Tools Customizer by running a REXX EXEC from the ISPF Command Shell panel.	"Starting Tools Customizer" on page 28	
Set up Tools Customizer user settings. If you are running Tools Customizer for the first time, you must modify several user settings to suit your environment. Otherwise, if the user settings that you have already established are still appropriate, skip this step.	"Modifying Tools Customizer user settings" on page 29	
Complete the steps in the appropriate customization roadmap based on the type of customization that you are performing.		

Task	Link to detailed instructions	Status
Customizing DB2 Sort for the first time Follow this roadmap if you do not have a customized version of DB2 Sort and you need to customize it for the first time.	“Roadmap: Customizing DB2 Sort for the first time” on page 33	
Customizing a different version of DB2 Sort Follow this roadmap if you have already customized a version of DB2 Sort and you want to use the same parameter values to customize a different version.	“Roadmap: Customizing a new version of DB2 Sort from a previous customization” on page 34	
Recustomizing DB2 Sort Follow this roadmap if you have a customized version of DB2 Sort but you want to change one or more parameter values.	“Roadmap: Recustomizing DB2 Sort” on page 35	
Accessing DB2 Sort libraries		
Create access to the SCNKLINK and SCNKLPA libraries.	“Accessing DB2 Sort libraries” on page 52	
Installation Verification Program		
Test the installation and customization using the Installation Verification Program (IVP).	Chapter 4, “Verifying installation and customization,” on page 53	

Worksheets: Gathering required data set names and parameter values

During the customization process, you will need to provide parameter values for DB2 Sort, and in some cases, for DB2.

Use the worksheets in this topic to record the appropriate parameter settings for your purposes, and then use these worksheets during the customization process. The worksheets are organized based on the order of the customization panels in Tools Customizer.

Tip: Print the following worksheets and refer to them during the customization process.

Discover customized product information

If you have already customized a version of DB2 Sort and you want to use the same parameter values to customize a different version, you can use the Discover EXEC in Tools Customizer. The Discover panel will prompt you for the name of the library from the previous version of DB2 Sort.

Parameter	Description	Your value
(Prior) SCNKLINK	The DB2 Sort SCNKLINK library from the previous version, such as CNK130.SCNKLINK	

Parameters: Data set names and installation options

Data set names for Tools Customizer

If you have not previously used Tools Customizer, you will need to

provide values for Tools Customizer data sets on the Tools Customizer Settings panel. In this worksheet, record the names of the Tools Customizer data sets.

Data set names for Tools Customizer

Data set name / hlq	Description	Special requirements	Your data set name / hlq
DB2TOOL.PRODUCT.CUST	Customization library qualifier	You must have write access to this data set qualifier.	
SCCQDENU	Metadata library for Tools Customizer	None.	
DISCOVER	Discover output data set	You must have write access to this data set.	
DATASTOR	Data store data set	You must have write access to this data set.	

Data set names for DB2 Sort

In this worksheet, record the names of the data sets for the DB2 Sort version that you are now customizing.

Data set names for DB2 Sort

Data set name	Description	Special requirements	Your data set name
SCNKLINK	DB2 Sort SCNKLINK library	This data set must be APF-authorized.	
SCNKLPA	DB2 Sort SCNKLPA library	If this data set is not in the LPA, it must be APF-authorized.	
SCNKUSER	DB2 Sort SCNKUSER library	You must have read access to this data set.	

DB2 Sort installation options

In this worksheet, record the values for the DB2 Sort installation options.

DB2 Sort installation options

Installation option	Delivered default value	Discovered?	Your value
Sort work DASD pool	SYSDA	Yes	
Sort work DASD type	3390	Yes	
Number of allocation retry attempts	5	Yes	
Minutes between retry attempts	3	Yes	
Sort work storage class	None	Yes	
Maximum number of sort works	32	Yes	
TUNEMSG	ON	Yes	
OPTMODE	ELAP	Yes	
PAGEMON	OFF	Yes	

IVP parameters

If you want to run the Installation Verification Program (IVP), you will need to supply values on the DB2 Parameters panels for the items in the following table. This task is optional.

IVP parameters

Parameter	Required?	Default value	Your value
Mode This parameter indicates the mode in which the DB2 subsystem is running.	Yes	NFM	
Level number This parameter indicates the version, release, and modification level of the DB2 subsystem.	Yes	blank	
Load library This parameter indicates the data set name of the DB2 load library.	Yes	DSN.SDSNLOAD	
Run library This parameter indicates the data set name of the DB2 run library.	Yes	DSN.RUNLIB.LOAD	
SYSAFF for DB2 utilities	No	blank	
Plan name for the DSNTEP2 utility	Yes	DSNTEP2	
Schema name	Yes	SYSTOOLS	
SQL authorization ID This ID is used in IVP JCL while creating or altering DB2 objects used by DB2 Sort.	Yes	blank	
Database name for the IVP	Yes	CNKDBNAM	
Storage group name for the IVP	Yes	CNKSTORG	
VCAT name for DB2 storage group for the IVP	Yes	CNKVCAT	
Unique identifier for DB2 utilities for the IVP	Yes	CNKUID	

IVP parameters

Parameter	Required?	Default value	Your value
Spare 4K buffer pool for IVP table space	Yes	BP48	
Current VPSIZE for 4K BP used by IVP table space	Yes	2000	
Spare 4K buffer pool for IVP index space	Yes	BP49	
Current VPSIZE for 4K BP used by IVP index space	Yes	2000	
Table space name for IVP table 1	Yes	CNKTBSP1	
Table space name for IVP table 2	Yes	CNKTBSP2	
Table name for IVP table 1	Yes	CNKTBLN1	
Table name for IVP table 2	Yes	CNKTBLN2	
Index name 1 for IVP table 1	Yes	CNKIDX11	
Index name 2 for IVP table 1	Yes	CNKIDX12	
Index name 3 for IVP table 1	Yes	CNKIDX13	
Index name 4 for IVP table 1	Yes	CNKIDX14	
Index name 5 for IVP table 1	Yes	CNKIDX15	
Index name 1 for IVP table 2	Yes	CNKIDX21	
Index name 2 for IVP table 2	Yes	CNKIDX22	
Index name 3 for IVP table 2	Yes	CNKIDX23	

Default installation options

The DB2 Sort default installation options include options for dynamic sort work allocation, optimization mode, sensitivity to paging activity, and tuning messages.

Dynamic sort work allocation options

The following options refer to values used for the dynamic allocation of sort work data sets. These data sets are de-allocated when the sort process has completed.

Sort work DASD pool

This value specifies the 1- to 8-character name of the group of devices from which DB2 Sort dynamically allocates the sort work data sets if SMS is not installed or active for temporary DASD work data sets. If SMS is active, you can specify a value for the sort work storage class option. The default value is SYSDA.

Restriction: You cannot specify VIO as a unit device.

Note: This value can be overridden at run time.

Sort work DASD type

This value specifies the type of DASD with the smallest track capacity that might be encountered at your installation by a dynamically allocated sort work data set. If, during dynamic allocation, DB2 Sort receives a device with a track capacity that is smaller than the one specified, it might ignore the specified device type. The default value is 3390.

The following table lists common DASD types and their track capacities:

Table 1. DASD types and their track capacities

DASD type	DASD track capacity
3380	47476
3390	56664
9345	46456

Number of allocation retry attempts

This value specifies the number of times that DB2 Sort automatically retries a dynamic allocation request in situations where sufficient disk space is not immediately available to satisfy that request. Insufficient disk space can result in a Sort Capacity Exceeded condition. The minimum value allowed is 0 and the maximum value is 16. The default value is 5.

Minutes between retry attempts

This value specifies the number of minutes to wait between dynamic allocation retry attempts. The minimum value allowed is 0 and the maximum value is 15. The default value is 3.

Sort work storage class

This value specifies the storage class that DB2 Sort uses to manage temporary work data sets for the dynamic allocation of sort work data sets if SMS is active. If this value is not specified in an SMS environment, DB2 Sort dynamically allocates devices from the generic class name that is specified for the sort work DASD pool. There is no default value for storage class.

Note: An installation-written automatic class selection (ACS) routine can override the value you specify.

Maximum number of sort works

This value specifies the maximum number of sort work data sets that can be used. DB2 Sort can dynamically allocate data sets to bring the total number of sort work data sets up to the number that you specify. Specify a

value from 32 through 255. DB2 Sort automatically determines the appropriate number of sort work files to use. The default value is 32.

Note: This value can be overridden at run time.

Optimization mode option

This value specifies the optimization mode for DB2 Sort.

OPTMODE

BALANCE

BALANCE provides the best mix of CPU and elapsed time performance by balancing utilization of central storage and disk space. Typically BALANCE uses less CPU time and more elapsed time than ELAP.

CPU CPU optimizes to minimize CPU time without any consideration for elapsed time, maximizing disk space usage.

ELAP ELAP optimizes to minimize elapsed time while significantly reducing CPU time, favoring central storage over disk space usage. The default value is ELAP.

Note 1: All of these options limit central storage use if auxiliary storage availability is low.

Note 2: The OPTMODE value can be overridden at run time.

Sensitivity to paging activity option

This value specifies the type of storage allocations that DB2 Sort chooses.

PAGEMON

ON Choose ON only if your system is highly sensitive to paging activity.

OFF Choose OFF if you use the defaults for DFSORT's DSPSIZE, HIPRMAX and MOSIZE installation options. The default value is OFF.

Note: Elapsed time may increase if PAGEMON is set to ON.

Tuning messages option

This value specifies whether DB2 Sort displays additional informational messages for tuning purposes. These tuning messages provide information about the product's resource utilization, such as memory usage, sort work DASD space, and sort work EXCPs.

TUNEMSG

ON Enables the generation of the additional tuning messages. The default value is ON.

OFF Suppresses the additional tuning messages.

Note: The TUNEMSG value can be overridden at run time.

Tools Customizer terminology and data sets

Before you use Tools Customizer, you should understand the Tools Customizer terminology and the data sets that Tools Customizer uses during customization.

Tools Customizer terminology

Tools Customizer uses several unique terms that you should be familiar with before you begin to use Tools Customizer.

Products and components

How an IBM Tool is packaged determines whether it is referred to as a product or as a component in the Tools Customizer documentation and interface. An IBM Tool that is ordered as a stand-alone entity (that is, not as part of a solution pack) is referred to as a product. An IBM Tool that is part of a solution pack is referred to as a component. Some IBM Tools are available in both formats; therefore, the same IBM Tool can be referred to as a product or as a component depending on how it is packaged.

DB2 entry

You can customize DB2 Sort on one or more DB2 entries. A DB2 entry can be any of the following items:

DB2 subsystem

A distinct instance of a relational database management system (RDBMS) that is not part of a data sharing group. An example of a DB2 subsystem name is DB01.

DB2 group attach name

The name that is used by the TSO/batch attachment, the call attachment facility (CAF), DL/I batch, utilities, and the Resource Recovery Services attachment facility (RRSAF) as a generic attachment name. An example of a group attach name is DSG1.

DB2 data sharing member

A DB2 subsystem that is assigned by the cross-system coupling facility (XCF) to a data sharing group. An example of a DB2 data sharing member name is DB02.

Tools Customizer maintains the following lists of DB2 entries:

Associated list

The list of DB2 entries that are associated with DB2 Sort. If the product to be customized requires DB2 entries, you can customize DB2 Sort only on DB2 entries that are in the associated list. When you customize DB2 Sort, this list is displayed in the DB2 Entries, Associations, and Parameter Status section of the Customizer Workplace panel.

You can add and copy DB2 entries to the associated list. When you add or copy DB2 entries to the associated list, the entries are associated with DB2 Sort.

Master list

The list of all DB2 entries that are defined but are not associated with DB2 Sort. Tools Customizer obtains information about these DB2 entries either from entries that were created manually or from the customizations of other products that were discovered. If you remove a DB2 entry from the associated list, the DB2 entry is

added to the master list. When you create a new DB2 entry, it is added to the master list, and when you associate the new entry with DB2 Sort, it is removed from the master list and added to the associated list. The master list is displayed on the Associate a DB2 Entry for Product panel.

If the associated list does not have the DB2 entries on which you want to customize DB2 Sort, you can associate existing entries from the master list to the associated list.

You can create new DB2 entries and copy existing entries to the master list.

High-level qualifier

The high-level qualifier is considered to be all of the qualifiers except the lowest level qualifier. A high-level qualifier includes a mid-level qualifier.

Product parameters

Parameters that are specific to DB2 Sort. These parameters are defined by DB2 Sort and are stored in a data member that is defined by DB2 Sort.

DB2 parameters

Parameters for a DB2 entry. These parameters are defined by Tools Customizer and are stored in a DB2 parameter data member.

Status type

Product, LPAR, and DB2 entry status type

After you specify the product that you want to customize, the product, the LPAR, and the DB2 entries have a status. The status is partly based on whether required parameters are defined. For some products, LPAR parameters or DB2 parameters might not be required. In these cases, the status is Not Required.

To customize DB2 Sort, all of the required parameters must be defined.

If required parameters for the the product parameters or DB2 parameters are not defined, the status of the parameters is Incomplete. Define values for parameters by manually editing them or by generating the customization jobs and specifying values for all of the required parameters that are displayed on the panels.

When values for all of the required parameters are defined, the status is Ready to Customize. Customization jobs can be generated only when all of the required parameters are defined and the status is Ready to Customize or Customized for the product parameters and DB2 parameters for the DB2 entries on which DB2 Sort will be customized.

The following table shows the meaning of the status types. Each status is defined differently for each type of parameter.

Table 2. Status types for the product, the LPAR, and the DB2 entries

Status	Product	LPAR	DB2 entries
Incomplete	The required product parameters are not defined, or the required product parameters are defined but LPAR parameters, DB2 parameters, or both are not defined.	The required parameters are not defined.	The required parameters are not defined.
Discovered	The product parameter definitions were discovered by using the product Discover EXEC.	N/A	N/A
Ready to Customize	The required product, LPAR, and DB2 parameters are defined, the status is Ready to Customize or Customized for the LPAR and at least one associated DB2 entry. You can generate the customization jobs.	The required LPAR parameters are defined or LPAR parameters are not required.	The required DB2 parameters are defined or DB2 parameters are not required.
Customized	The jobs are customized on the local LPAR.	The jobs are customized for the product or for all of the associated DB2 entries on the local LPAR.	The jobs are customized for the DB2 entry.
Errors in Customization	N/A	N/A	Errors occurred while the customization jobs were being generated.
Not Required	N/A	LPAR parameters are not required.	DB2 parameters are not required.

Related tasks:

“Creating and associating DB2 entries” on page 40

DB2 entries are optional for DB2 Sort. You can create new DB2 entries and associate them with DB2 Sort.

“Copying DB2 entries” on page 47

You can copy associated and not associated DB2 entries to other DB2 entries or to new DB2 entries.

“Removing DB2 entries” on page 49

You can remove DB2 entries from the associated list.

Data sets that Tools Customizer uses during customization

Tools Customizer uses several unique data sets during the customization process. Familiarize yourself with these data sets before you begin to use Tools Customizer.

Several different data sets are required to customize DB2 Sort with Tools Customizer. These data sets are supplied by DB2 Sort, supplied by Tools Customizer, or allocated by Tools Customizer.

DB2 Sort provides the following data sets:

Metadata library

Contains the metadata for the product to be customized. Tools Customizer uses the metadata to determine which tasks, steps, and parameters to display on the Product Parameters panel, the LPAR Parameters panel, and the DB2 Parameters panel. This data set also contains the templates that Tools Customizer uses to generate the customization jobs.

The metadata library naming convention is *high_level_qualifier*.SCNKDENU, where *high_level_qualifier* is all of the segments of the data set name except the lowest-level qualifier.

You specify the metadata library on the Specify the Metadata Library panel. READ access to this data set is required.

Discover EXEC library

Contains the DB2 Sort Discover EXEC. When you customize DB2 Sort, you can use the Discover EXEC to automatically retrieve and store product information, such as parameter values from an already customized product. Tools Customizer saves the discovered information in the data store.

The default name of the data set is the high-level qualifier for the metadata library plus a lowest-level qualifier. For DB2 Sort, the lowest-level qualifier is SCNKDENU. You can change the default value on the Discover Customized Product Information panel. EXECUTE access to this data set is required.

Tools Customizer provides the following data sets:

Tools Customizer metadata library

Contains the metadata for the DB2 and LPAR parameters that are required to customize DB2 Sort. Tools Customizer uses the metadata to determine which parameters to display on the DB2 Parameters panel and the LPAR Parameters panel. In addition, Tools Customizer uses information in the metadata library to determine whether additional DB2 and LPAR parameters need to be displayed on these panels. As you customize different products, different DB2 and LPAR parameters might need to be defined.

The default name of the data set is DB2TOOL.CCQ110.SCCQDENU. You can change the default value on the Tools Customizer Settings panel. READ access to this data set is required.

Tools Customizer table library

Stores information about jobs that are customized. Job information that is stored includes a description of the job, its member name and template name, the SSID, group attach name, and when the job was generated.

The default name of the data set is DB2TOOL.CCQ110.SCCQTENU. WRITE access to this data set is required.

Tools Customizer requires that the following data sets exist during the customization process. If the data sets do not exist, Tools Customizer automatically allocates them.

Discover output data set

Contains the output that is generated when you run the DB2 Sort Discover EXEC. The DB2 Sort Discover EXEC retrieves the metadata and values for the parameters from a previous customization of DB2 Sort.

The default name of the data set is DB2TOOL.CCQ110.DISCOVER. You can change the default value on the Tools Customizer Settings panel or the Discover Customized Product Information panel. WRITE access to this data set is required.

Data store data set

Contains product, LPAR, and DB2 parameter values, and DB2 entry associations. Tools Customizer uses this data set to permanently store all information that is acquired about the product, DB2 subsystems or data sharing groups, and LPAR when you customize products on the local LPAR.

The default name of the data set is DB2TOOL.CCQ110.DATASTOR. You can change the default value on the Tools Customizer Settings panel. WRITE access to this data set is required.

Customization library

Contains the customization jobs that Tools Customizer generates for DB2 Sort.

Tools Customizer checks whether a customization library name was specified for more than one instance of the same version of the same product. If the same customization library name is specified for more than one product of the same version, the CCQD123E message is issued to prevent you from overwriting previously generated customization jobs. Ensure that you specify unique qualifier for the customization library for each instance of the product.

To customize DB2 Sort, submit the members of the data set in the order in which they are displayed on the Finish Product Customization panel.

The data set naming convention is *hlq.\$LPAR_name\$.xyzvrm*, where:

- *hlq* is the value of the **Customization library qualifier** field on the Tools Customizer Settings panel (CCQPSET)
- *LPAR_name* is the four-character LPAR name
- *xyzvrm* is the three-letter product identifier with the version, release, and modification level

For example, the data set name might be DB2TOOL.PRODUCT.CUST.\$MVS1\$.XYZ410.

WRITE access to this data set is required.

Tools Customizer allocates the data sets for the discover output, the data store, and the customization library with the attributes that are shown in the following table:

Table 3. Data set attributes for allocating the Discover output, data store, and customization library data sets

Data set	Organization	Record format	Record length	Block size	Data set name type
Discover output data set	PO	Variable block	16383	32760	LIBRARY

Table 3. Data set attributes for allocating the Discover output, data store, and customization library data sets (continued)

Data set	Organization	Record format	Record length	Block size	Data set name type
Data store data set	PO	Variable block	16383	32760	LIBRARY
Product customization library	PO	Fixed block	80	32720	LIBRARY

Restrictions:

- Multiple users cannot simultaneously share the discover output data set, data store data set, Tools Customizer metadata library, and metadata library.
- You cannot share the data store data set across multiple LPARs with shared DASD or copy the data store data set to another LPAR. Tools Customizer creates many cross-references between product and DB2 associations. Therefore, if you share or copy the data store data set, member names that are empty or that do not exist might be generated.

Runtime considerations

To improve performance, you might want to modify the **REGION** parameter in your DB2 utility job step. You can also override some of the installation options at run time.

The REGION parameter

Because DB2 Sort manages memory allocation dynamically, setting REGION=0M is recommended for best performance.

Default installation options that you can override at run time

Although most applications can use the default installation options, some applications might require different options. You can override installation options at run time by passing parameters in the DB2SORTI DD data set.

Important: This data set must have a fixed-length format with 80-byte records.

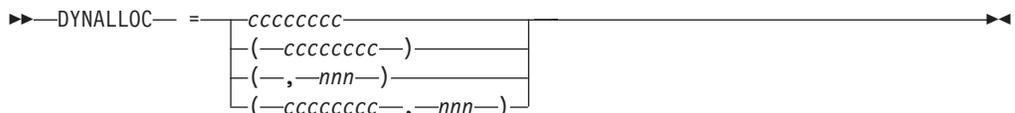
You can specify the **DYNALLOC**, **OPTMODE**, and **TUNEMSG** parameters at run time.

The parameters can start in any column but must be completed on the current record.

DYNALLOC

This parameter specifies values for dynamic sort work allocation.

Syntax



Example

```
//DB2SORTI DD *  
DYNALLOCL(,50),OPTMODE=ELAP
```

This example of a runtime override accomplishes the following goals:

- Specifies 50 for the maximum number of sort work data sets that can be used
- Sets the optimization mode to ELAP

Chapter 3. Activating and customizing DB2 Sort

You must activate DB2 Sort after it is installed on your system. After activating DB2 Sort, you can customize it by changing the default installation options with Tools Customizer. After customizing DB2 Sort, you must create access to the DB2 Sort libraries, then you can use Installation Verification Program to test your settings before moving to a production environment.

Note: For the best resource utilization, set the **UTSORTAL** parameter to YES. This will enable the best determination of the number of parallel sort tasks depending on the number of data sets allocated by DB2.

For information about setting the **UTSORTAL** parameter, see one of the following documents:

- *DB2 11 for z/OS Installation and Migration Guide (GC19-4056)*
- *DB2 10 for z/OS Installation and Migration Guide (GC19-2974)*

Topics:

- “Activating DB2 Sort”
- “Starting Tools Customizer” on page 28
- “Customizing DB2 Sort” on page 33
- “Accessing DB2 Sort libraries” on page 52

Activating DB2 Sort

Before customizing or using DB2 Sort, you must activate it or verify that it has been activated.

About this task

You may have to activate DB2 Sort for use in DB2 for z/OS if it has been migrated from a release earlier than DB2 10 for z/OS.

Important: If you are using a newly installed DB2 10 for z/OS or DB2 11 for z/OS, DB2 Sort is automatically activated, so no action is required.

To activate DB2 Sort, use the following procedure.

Procedure

1. Apply the following program fix APAR PI19784: DB2 SORT FOR Z/OS PRECONDITIONING.
2. Set the DB2 subsystem parameter **DB2SORT** to ENABLE.

For information about setting the **DB2SORT** parameter, see one of the following documents:

- *DB2 11 for z/OS Installation and Migration Guide (GC19-4056)*
- *DB2 10 for z/OS Installation and Migration Guide (GC19-2974)*

What to do next

1. Stop and restart the DB2 subsystem.
2. Customize, test or begin using DB2 Sort.

Starting and preparing Tools Customizer for use

Use the provided REXX EXEC to start Tools Customizer. The first time that you use Tools Customizer, you must modify the settings that Tools Customizer uses to customize DB2 Sort.

Starting Tools Customizer

Start Tools Customizer by running a REXX EXEC from the ISPF Command Shell panel.

Before you begin

Tools Customizer must be SMP/E installed. You must know the high-level qualifier of where the Tools Customizer libraries reside. The high-level qualifier is considered to be all the segments of the data set name except the lowest-level qualifier, which is SCCQEXEC.

Attention: Ensure that Tools Customizer load libraries are not APF authorized. APF authorizing Tools Customizer libraries results in an abend.

About this task

To run the REXX EXEC, you must either change the placeholder in the EXEC for the high-level qualifier of the Tools Customizer EXEC library or pass the high-level qualifier as a parameter when you run the EXEC. The REXX EXEC is in the CCQTCZ member of the EXEC library.

Procedure

1. Optional: Change the placeholder for the high-level qualifier in the REXX EXEC:
 - a. Find the EXEC library data set for Tools Customizer. The name of the data set is *high_level_qualifier.SCCQEXEC*.
 - b. Edit data set member CCQTCZ and replace the <TCZ HLQ> string with the high-level qualifier of the EXEC library data set. For example, if the name of the Tools Customizer EXEC library is CCQTCZ.USABSAND.SCCQEXEC, replace <TCZ HLQ> with CCQTCZ.USABSAND.

You have to change the placeholder for the high-level qualifier only once. When you run the REXX EXEC, you do not have to pass the high-level qualifier as a parameter.

2. Run the REXX EXEC (CCQTCZ):
 - a. From the ISPF Primary Option Menu, select option 6. The ISPF Command Shell panel is displayed.
 - b. Specify the EX command to run the REXX EXEC. For example, if the Tools Customizer EXEC library is CCQTCZ.USABSAND.SCCQEXEC and you changed the placeholder for the high-level qualifier in the REXX EXEC, specify: EX 'CCQTCZ.USABSAND.SCCQEXEC(CCQTCZ)'

If you did not change the placeholder for the high-level qualifier in the REXX EXEC, specify: EX 'CCQTCZ.USABSAND.SCCQEXEC(CCQTCZ)'
'CCQTCZ.USABSAND'

Results

The IBM Customizer Tools for z/OS main menu panel is displayed.

What to do next

If you are running Tools Customizer for the first time, you must modify the Tools Customizer user settings. If you have already set the Tools Customizer user settings, either customize or recustomize DB2 Sort.

Modifying Tools Customizer user settings

Before you can customize DB2 Sort with Tools Customizer, you must review the settings that Tools Customizer uses. You might have to change the default values to suit your environment. In most cases, you can change the Tools Customizer values at any time. For example, after you have customized DB2 Sort and are customizing a different product or solution pack, you might have to change the settings.

Procedure

1. On the IBM Tools Customizer for z/OS main panel (CCQPHME), specify option 0, **User settings for Tools Customizer**. The Tools Customizer Settings panel (CCQPSET) is displayed, as shown in the following figure:

```
CCQPSET          Tools Customizer Settings          14:03:51
Command ==>>
Enter the settings for customizing a product or press End to save and exit.

Commands: SAVE - Save user settings

Product Customization Settings
Customization library qualifier . . DB2TOOL.PRODUCT.CUST
Use DB2 group attach . . . . . YES (YES/NO)

Tools Customizer Library Settings
Metadata library . . . . . DB2TOOL.CCQ110.SCCQDENU
Discover output data set . DB2TOOL.CCQ110.DISCOVER
Data store data set . . . DB2TOOL.CCQ110.DATASTOR

User Job Card Settings for Customization Jobs
==>> //          JOB
==>>
==>>
==>>
==>>
```

Figure 2. The Tools Customizer Settings panel (CCQPSET)

2. Review the values for the following required fields. Use the default value or specify your own value. You must have appropriate read and write access to the data sets that are specified.

Customization library qualifier

The high-level qualifier that is used as the prefix for the customization library. The customization library is a data set in which the generated jobs to customize DB2 Sort are stored. Write access to this qualifier is required.

For each product to be customized, the first value that is specified for the qualifier is always used, even if you change it after you have generated the customization jobs. For example, if you customize a

product and then specify a new qualifier for recustomization, although the new qualifier is saved and displayed, the original value is used.

To maintain multiple instances of Tools Customizer, specify a unique customization library qualifier for each instance of Tools Customizer. Data set names that exceed 42 characters must be enclosed in single quotation marks (').

Use DB2 group attach

Determines the value that is used in the CONNECT statements in the generated customization jobs. Specify YES for data sharing environments, which causes the group attach name to be used. Specifying NO, in most cases, causes the SSID to be used in the DB2 CONNECT statement.

Important: This field has no effect when you are customizing a product on a DB2 subsystem that is not a member of a data sharing group. In this case, the DB2 subsystem ID (SSID) is always used in the CONNECT statements in the generated customization jobs.

When you are customizing a product on a DB2 subsystem that is a member of a data sharing group, how the DB2 subsystem is defined and the value of the **Use DB2 group attach** field determines the value that is used in the CONNECT statements in the generated jobs. The following table shows whether the SSID or the group attach name is used:

Table 4. The effect of the value of the Use DB2 group attach field in a data sharing environment

DB2 subsystem definition	Value of the Use DB2 group attach field	Value that is used in the CONNECT statements
The DB2 subsystem is defined with an SSID.	Yes	Group attach name
	No	SSID ¹
The DB2 subsystem is not defined with an SSID.	Yes or No	Group attach name

Note 1: If you generate jobs for multiple DB2 subsystems that are defined with an SSID and belong to the same data sharing group, the SSID of the first DB2 subsystem that is selected is used.

For example, assume that on the Customizer Workplace panel, you generated jobs for the following DB2 subsystems:

- V91C, which is a stand-alone DB2 subsystem
- V91A, which is a DB2 subsystem that is a member of data sharing group DSG1
- A DB2 subsystem that was not defined with an SSID that is a member of data sharing group DSGA

The following figure shows how these DB2 entries might be listed on the Customizer Workplace panel:

```

Associated DB2 Entries and Parameter Status
Line commands: G - Generate jobs  E - Edit  B - Browse  C - Copy  R - Remove
Cmd SSID GrpAttch Lvl Mode User ID Date      Status      Message
V91C  --          910 NFM  SYSADM 2010/11/09 Ready to Customize
V91A  DSG1       910 NFM  SYSADM 2010/11/09 Ready to Customize
--    DSGA       910 NFM  SYSADM 2010/11/09 Ready to Customize
----- End of DB2 entries -----

```

The following table shows which values are used in the CONNECT statements in the generated jobs, based on the value of the **Use DB2 group attach** field.

Table 5. Value that is used in the CONNECT statements in the generated jobs

SSID	GrpAttch	Value of the Use DB2 group attach field	Value that is used in the CONNECT statements
V91C	--	Yes	SSID
		No	SSID
V91A	DSG1	Yes	Group attach name
		No	SSID
--	DSGA	Yes	Group attach name
		No	Group attach name

Tools Customizer metadata library

The name of the data set that contains the metadata that is used to display the DB2 parameters. The parameters that are displayed on the DB2 Parameters panel depend on the parameters that you define and the tasks and steps that you select on the Product Parameters panel for the product that you are customizing. For example, the DB2 parameters that are required, based on the selected tasks and steps, are displayed on the DB2 Parameters panel, and you can edit them. If they are not required, they are not displayed. Read access to this data set is required. Data set names that exceed 42 characters must be enclosed in single quotation marks (').

Discover output data set

The name of the data set in which the output from the DB2 Sort Discover EXEC is stored. Each product has its own Discover EXEC. The Discover EXEC retrieves the product and DB2 parameters from a previously customized product. Write access to this data set is required. Data set names that exceed 42 characters must be enclosed in single quotation marks (').

Data store data set

The name of the data set where Tools Customizer stores information about product and DB2 parameter values. Information about which products are associated with which DB2 entries (DB2 subsystems, DB2 group attach names, and DB2 data sharing members) is also stored in this data set. Data set names that exceed 42 characters must be enclosed in single quotation marks ('). The specified data store data set can be used with only one invocation of Tools Customizer at a time. Data set names that exceed 42 characters must be enclosed in single quotation marks (').

User job card settings for customization jobs

The job card information to be inserted into the generated jobs for

customizing a product. The default value is the job statement information from the ISPF Batch Selection panel.

The first line of the job card automatically begins with the following information:

```
//          JOB
```

where characters 3 - 10 are reserved by Tools Customizer for the job name and includes a blank space after JOB. This name cannot be edited. Information that you specify on the first line of the job card cannot exceed 57 characters. This character limit includes a continuation character. All other lines of the job card cannot exceed 72 characters.

3. Press End to save and exit. If the Discover output data set and the data store data set that you specified do not exist, Tools Customizer creates them.

Important: If the ISPF sessions unexpectedly ends before you exit Tools Customizer, the fields on the Tools Customizer Settings panel (CCQPSET) will be repopulated with default values, and you will be required to review them or specify new values again.

Results

The values are saved, and the IBM Tools Customizer for z/OS main menu panel (CCQPHME) is displayed again.

What to do next

You are ready to customize or recustomize DB2 Sort or to change parameter settings.

Related concepts:

“Customizing DB2 Sort” on page 33

Using Tools Customizer to customize DB2 Sort consists of identifying the product to customize; defining any required DB2 Sort and DB2 parameters; generating the customization jobs; and submitting the jobs.

Hiding and displaying panel text

After you are familiar with Tools Customizer, you might want to hide the instructions and some of the basic descriptions that are displayed by default on Tools Customizer panels.

About this task

By using the OPTIONS command, you can choose to show or hide the following information on Tools Customizer panels:

- The instructions on all panels
- The Product to Customize section on the Customizer Workplace panel (CCQPWRK)
- The Usage Notes section on the Product Parameters panel (CCQPPRD), the LPAR Parameters panel (CCQPLPR), and the DB2 Parameters panel (CCQPDB2).

By hiding this information, more data can be displayed on the panels. Later, you can redisplay this information also by using the OPTIONS command.

Procedure

1. On any Tools Customizer panel, issue the OPTIONS command. The Panel Display Options panel (CCQPOPT) is displayed, as shown in the following figure. By default, all options are preselected with a slash (/) to be shown.

```
CCQPOPT                      Panel Display Options

Select panel display options and press Enter. To cancel, press End.

Panel Display Options
Specify a slash (/) to select options.
/ Show the panel instructions
/ Show the Product to Customize section
/ Show the Usage Notes section
Command ==>
```

Figure 3. The Panel Display Options panel (CCQPOPT)

2. To hide any of the options, remove the slash, and press Enter.

Customizing DB2 Sort

Using Tools Customizer to customize DB2 Sort consists of identifying the product to customize; defining any required DB2 Sort and DB2 parameters; generating the customization jobs; and submitting the jobs.

Customization roadmaps describe the steps that you must complete to customize DB2 Sort. Separate roadmaps are provided for the three most common types of customizations.

Use the following table to determine which roadmap corresponds to your environment.

Table 6. Customization roadmaps

Environment description	Roadmap
You do not have a customized version of DB2 Sort, and you need to customize it for the first time.	"Roadmap: Customizing DB2 Sort for the first time"
You have already customized a version of DB2 Sort, and you want to use the same parameter values to customize a different version.	"Roadmap: Customizing a new version of DB2 Sort from a previous customization" on page 34
You have a customized version of of DB2 Sort, but you want to change one or more parameter values.	"Roadmap: Recustomizing DB2 Sort" on page 35

Roadmap: Customizing DB2 Sort for the first time

This roadmap lists and describes the steps that are required to customize DB2 Sort for the first time.

If you are customizing a previous version of DB2 Sort, see "Roadmap: Customizing a new version of DB2 Sort from a previous customization" on page 34.

Before you complete these steps, ensure that the following prerequisites have been met:

- All of the product customization steps that must be done before Tools Customizer is started are complete.
- Tools Customizer is started.
- The Tools Customizer settings have been reviewed or modified, and saved.

Complete the steps in the following table to customize DB2 Sort for the first time.

Table 7. Steps for customizing DB2 Sort for the first time

Step	Description	Instructions
1	Specify the product metadata library for the product that you want to customize. The name of this library is <i>hlq.SCNKDENU</i> .	"Specifying the metadata library for the product to customize" on page 36
2	Create new DB2 entries and associate them with DB2 Sort.	"Creating and associating DB2 entries" on page 40
3	Define the required parameters.	"Defining parameters" on page 41
4	Generate the customization jobs for the product or for the DB2 entries on which DB2 Sort is ready to be customized.	"Generating customization jobs" on page 45
5	Submit the generated customization jobs.	"Submitting customization jobs" on page 45

The following table lists some of the common administrative tasks that you might need to do during the customization process.

Table 8. Administrative tasks

Description	Instructions
Browse the different types of parameters.	"Browsing parameters" on page 47
Copy an existing DB2 entry to the list of DB2 entries on which DB2 Sort can be customized.	"Copying DB2 entries" on page 47
Remove one or more DB2 entries from the associated list.	"Removing DB2 entries" on page 49
Delete one or more DB2 entries from the master list.	"Deleting DB2 entries" on page 49
Display a list of customization jobs that have been previously generated.	"Displaying customization jobs" on page 50
Maintain the customization jobs in the customization library.	"Maintaining customization jobs" on page 50

Roadmap: Customizing a new version of DB2 Sort from a previous customization

This roadmap lists and describes the steps for customizing a new version of DB2 Sort based on the existing customization values of a previous version of the same product.

Use this roadmap even if the previous version of DB2 Sort was not customized by using Tools Customizer.

Before you complete these steps, ensure that the following prerequisites have been met:

- All of the product customization steps that must be done before Tools Customizer is started are complete.
- Tools Customizer is started.
- The Tools Customizer settings have been reviewed or modified, and saved.

Complete the steps in the following table to customize a new version of DB2 Sort from a previous customization.

Table 9. Steps for customizing a new version of DB2 Sort from a previous customization

Step	Description	Instructions
1	Specify the product metadata library for the product that you want to customize. The name of this library is <i>hlq.SCNKDENU</i> .	"Specifying the metadata library for the product to customize" on page 36
2	Use the DB2 Sort Discover EXEC to discover information about the version of DB2 Sort that you previously customized manually.	"Discovering DB2 Sort information automatically" on page 38
3	Define the required parameters.	"Defining parameters" on page 41
4	Generate the customization jobs for the product or for the DB2 entries on which DB2 Sort is ready to be customized.	"Generating customization jobs" on page 45
5	Submit the generated customization jobs.	"Submitting customization jobs" on page 45

The following table lists some of the common administrative tasks that you might need to do during the customization process.

Table 10. Administrative tasks

Description	Instructions
Browse the different types of parameters.	"Browsing parameters" on page 47
Copy an existing DB2 entry to the list of DB2 entries on which DB2 Sort can be customized.	"Copying DB2 entries" on page 47
Remove one or more DB2 entries from the associated list.	"Removing DB2 entries" on page 49
Delete one or more DB2 entries from the master list.	"Deleting DB2 entries" on page 49
Display a list of customization jobs that have been previously generated.	"Displaying customization jobs" on page 50
Maintain the customization jobs in the customization library.	"Maintaining customization jobs" on page 50

Roadmap: Recustomizing DB2 Sort

This roadmap lists and describes the steps to change parameter values and regenerate customization jobs for DB2 Sort after you have customized it for the first time.

The new customization jobs will replace the customization jobs that were previously generated and stored in the customization library. Part of the recustomization process includes selecting or deselecting optional tasks or steps, changing the definitions of parameters that have already been defined, or both. Use the method in this roadmap instead of deleting customization jobs from the customization library.

Before you complete these steps, ensure that the following prerequisites have been met:

- All of the product customization steps that must be done before Tools Customizer is started are complete.
- Tools Customizer is started.

Complete the steps in the following table to recustomize DB2 Sort.

Table 11. Required steps for recustomizing DB2 Sort

Step	Description	Instructions
1	Specify the product metadata library for the product that you want to recustomize. The name of this library is <i>hlq.SCNKDENU</i> .	"Specifying the metadata library for the product to customize"
2	Edit the specific tasks, steps, or parameters that need to be changed.	<ul style="list-style-type: none"> • "Defining DB2 Sort parameters" on page 42 • "Defining DB2 parameters" on page 43
3	Generate the customization jobs for the product or for the DB2 entries on which DB2 Sort is ready to be customized.	"Generating customization jobs" on page 45
4	Submit the new generated customization jobs.	"Submitting customization jobs" on page 45

The following table lists some of the common administrative tasks that you might need to do during the customization process.

Table 12. Administrative tasks

Description	Instructions
Browse the different types of parameters.	"Browsing parameters" on page 47
Copy an existing DB2 entry to the list of DB2 entries on which DB2 Sort can be customized.	"Copying DB2 entries" on page 47
Remove one or more DB2 entries from the associated list.	"Removing DB2 entries" on page 49
Delete one or more DB2 entries from the master list.	"Deleting DB2 entries" on page 49
Display a list of customization jobs that have been previously generated.	"Displaying customization jobs" on page 50
Maintain the customization jobs in the customization library.	"Maintaining customization jobs" on page 50

Specifying the metadata library for the product to customize

You must specify a metadata library for the product that you want to customize.

About this task

The product metadata library contains the information that determines which tasks, steps, and parameters are required to customize DB2 Sort. This information controls what is displayed on the Product Parameters panel and the DB2 Parameters panel.

After DB2 Sort has been SMP/E installed, the default name of the product metadata library is *high_level_qualifier*.SCNKDENU, where *high_level_qualifier* is all of the segments of the data set name except the lowest-level qualifier.

Procedure

1. Specify option 1 on the Tools Customizer for z/OS panel. The Specify the Metadata Library panel is displayed. This panel contains a list of the product metadata libraries that you specified most recently. If you are using Tools Customizer for the first time, this list is empty, as shown in the following figure:

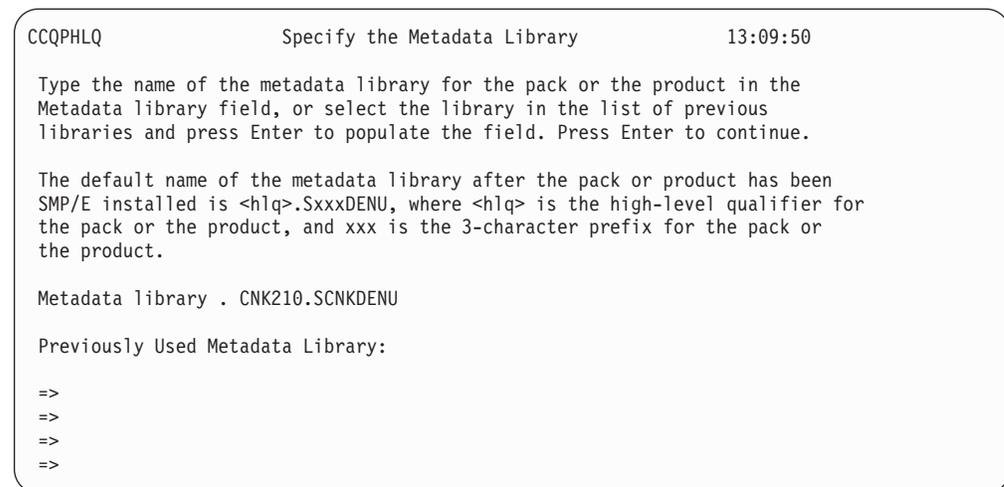


Figure 4. The Specify the Metadata Library panel

2. Use one of the following methods to specify the product metadata library:
 - Type the name of a fully qualified partitioned data set (PDS) or an extended partitioned data set (PDSE) in the **Metadata library** field. Double quotation marks (") cannot be used around the name. Single quotation marks (') can be used but are not required. If you are customizing DB2 Sort for the first time, you must use this method.
 - Place the cursor on the library name in the Recent Metadata Libraries list, and press Enter.

Results

If you are customizing DB2 Sort for the first time, the Run Discover EXEC panel is displayed. Otherwise, the Customizer Workplace panel is displayed.

What to do next

- Complete the steps that correspond to your environment:

Customizing DB2 Sort for the first time

Do not run the DB2 Sort Discover EXEC. Press End. The Customizer Workplace panel is displayed. If your environment requires associated

DB2 entries, ensure that they are created and associated. If your environment does not require associated DB2 entries, skip this step, and edit DB2 Sort parameters.

Customizing DB2 Sort from a previous or current customization

Press Enter to run the DB2 Sort Discover EXEC. The Discover Customized Product Information panel is displayed. Specify the required information for running the EXEC.

Discovering DB2 Sort information automatically

You can use the DB2 Sort Discover EXEC to discover information from a previous or current customization of .

About this task

Tip: Using the DB2 Sort Discover EXEC to discover information from a previous or current customization saves time and reduces errors that can occur when parameters are specified manually.

DB2 Sort provides the Discover EXEC that you will run. Therefore, the information that can be discovered depends on DB2 Sort.

Parameter values that are discovered and parameter values that are specified manually are saved in the data store. If parameter values for the product that you want to customize exist in the data store, Tools Customizer issues a warning before existing values are replaced.

Procedure

1. On the Customizer Workplace panel, issue the DISCOVER command. If you chose to run the DB2 Sort Discover EXEC on the pop-up panel after you specified the product to customize, skip this step.

Tip: You can run any Tools Customizer primary command by using either of the following methods:

- Place the cursor on the name of the primary command, and press Enter.
- Type the primary command name in the command line, and press Enter.

The Discover Customized Product Information panel is displayed, as shown in the following figure:

```

CCQPDSC          Discover Customized Product Information          10:49:05
Command ==>>>          Scroll ==>> CSR

For the product you are customizing, the Discover EXEC retrieves product
information from an already customized product. Specify the required
information. To save your information and run the Discover EXEC, issue the RUN
command. To save your information and stay on this panel, issue the SAVE
command. To verify the syntax of your information without saving it, press
Enter. To save and exit, press End.

Commands: RUN  SAVE

Discover EXEC for Extracting Information from an Already Customized product
Discover EXEC library . . . CNK210.SCNKDENU
Discover EXEC name . . . . : CNKDISCV
Discover output data set . . DB2TOOL.CCQ110.DISCOVER

Information for Discover EXEC
Prior DB2 Sort SCNKLINK library
CNK130.SCNKLINK

```

Figure 5. The Discover Customized Product Information panel

2. Either accept the default values for the following input fields that Tools Customizer generates, or replace the default values with your own values:
 - Discover EXEC library**
The fully qualified data set name that contains the DB2 Sort Discover EXEC.
 - Discover EXEC name**
The name of the DB2 Sort Discover EXEC.
 - Discover output data set**
The fully qualified data set where output from the DB2 Sort Discover EXEC is stored.
3. Either accept or change the default values in the **Information for Discover EXEC** fields. These fields are generated by DB2 Sort. They show the information that is required to run the DB2 Sort Discover EXEC.
4. Issue the RUN command to run the DB2 Sort Discover EXEC. Alternatively, save your information without running the DB2 Sort Discover EXEC by issuing the SAVE command. If you issue the RUN command to run the DB2 Sort Discover EXEC, the parameter information is discovered for DB2 Sort, and the Customizer Workplace panel is displayed.

Results

The discovered parameter values for DB2 Sort replace any existing values.

What to do next

The next step depends on your environment:

- If DB2 entries were not discovered, or if you need to customize DB2 Sort on new DB2 entries, create and associate the entries.
- If DB2 entries were discovered and you want to customize DB2 Sort on only these entries, define the parameters.

Related tasks:

“Creating and associating DB2 entries” on page 40

DB2 entries are optional for DB2 Sort. You can create new DB2 entries and

associate them with DB2 Sort.

“Defining parameters” on page 41

To customize DB2 Sort, you must define DB2 Sort parameters and DB2 parameters, if your customization requires DB2 entries.

Creating and associating DB2 entries

DB2 entries are optional for DB2 Sort. You can create new DB2 entries and associate them with DB2 Sort.

About this task

The list of associated DB2 entries is on the Customizer Workplace panel.

Procedure

1. Issue the ASSOCIATE command on the Customizer Workplace panel. The Associate DB2 Entry for Product panel is displayed, as shown in the following figure:

```
CCQPDAD          Associate DB2 Entry for Product          11:23:03
Command ==>>>          Scroll ==>> CSR

Select any of the following DB2 entries to add them to the Customizer
Workplace panel. You use the Customizer Workplace panel to choose the DB2
subsystems, data sharing members, and group attach names on which to
customize the product.

Commands: CREATE - Create new DB2 entries

DB2 Entries
Line commands: A - Associate C - Copy D - Delete
Cmd SSID GrpAttch
----- End of DB2 entries -----
```

Figure 6. The Associate DB2 Entry for Product panel

2. Create DB2 entries. If you need to associate DB2 entries that are already in the master list, skip this step and go to step 3.
 - a. Issue the CREATE command. The Create DB2 Entries panel is displayed, as shown in the following figure:

```
TCUSTOMZR          Create DB2 Entries          Row 1 to 1 of 1
Command ==>>>          Scroll ==>> PAGE

Specify the SSID, the group attach name, or both in the appropriate columns
for each new DB2 entry and press Enter. To create additional entries, issue
the Inn line command, where nn is the number of entries to be inserted, and
press Enter. To cancel, press End.

New DB2 Entries
Line commands: I - Insert into list R - Remove from list
Cmd SSID GrpAttch Message
```

Figure 7. The Create a DB2 Entry panel

- b. In the appropriate columns, specify a DB2 subsystem ID, DB2 group attach name, or DB2 data sharing member name for the DB2 entry that you want to create, and press Enter. Valid values are 1 - 4 characters. You can use symbolic characters. You cannot use blanks.

Tips:

- To insert multiple DB2 entries, specify the *Inn* line command, where *nn* is the number of DB2 entries to be inserted.
- You will define specific parameters for these new DB2 entries, such as parameters that define a subsystem as primary, on the DB2 Parameters panel. This panel is displayed after you select these new DB2 entries and issue the line command to generate the jobs, after you issue the primary command to generate the jobs for all associated DB2 entries, or when you manually edit the DB2 parameters.

The Associate DB2 Entry for Product panel is displayed, and the new DB2 entry is displayed in the master list, as shown in the following figure:

```

CCQPDAD                Associate DB2 Entry for Product                Scroll ==> CSR
Command ==>

Select any of the following DB2 entries to add them to the Customizer
Workplace panel. You use the Customizer Workplace panel to choose the DB2
subsystems, data sharing members, and group attach names on which to
customize the product.

Commands:  CREATE - Create new DB2 entries

DB2 Entries
Line commands: A - Associate  C - Copy  D - Delete
Cmd SSID GrpAttch
DB02  --
----- End of DB2 entries -----

```

Figure 8. The Associate DB2 Entry for Product panel with a new DB2 entry in the master list

- c. Repeat steps b and c for each DB2 entry that you want to create.
 - d. When you have created all the DB2 entries, associate them with DB2 Sort, or press End to display the Customizer Workplace panel.
3. Associate DB2 entries.
 - a. Specify A against one or more DB2 entries in the master list, and press Enter to associate them with DB2 Sort.

Results

The Customizer Workplace panel is displayed with the associated DB2 entries displayed in the associated list.

What to do next

Define the parameters.

Related concepts:

“Tools Customizer terminology” on page 18

Tools Customizer uses several unique terms that you should be familiar with before you begin to use Tools Customizer.

Defining parameters

To customize DB2 Sort, you must define DB2 Sort parameters and DB2 parameters, if your customization requires DB2 entries.

About this task

You must define the DB2 Sort parameters first for the following reasons:

- If you ran the DB2 Sort Discover EXEC, you must review the values that were discovered.
- If you select optional tasks and steps on the Product Parameters panel that affect the DB2 entry on which you will customize DB2 Sort, additional parameters might be displayed on the DB2 Parameters panel.
- If other steps must be completed in a specific sequence, customization notes on the Product Parameters panel will display the correct sequence.

Defining DB2 Sort parameters

DB2 Sort parameters are specific to DB2 Sort.

About this task

If you ran the DB2 Sort Discover EXEC, you must review the parameters that were discovered.

Procedure

1. Specify E next to the **Product parameters** field on the Customizer Workplace panel, and press Enter. The Product Parameters panel is displayed, as shown in the following figure. If other steps must be completed in a specific sequence before you define the DB2 Sort parameters, a note labeled **Important** will display the correct sequence on this panel.

```

CCQPPRD          Product Parameters: DB2 Sort          11:28:06
Command ==>>>          Scroll ==>> CSR

Complete the following tasks to customize the products. The required tasks
and steps are preselected. Ensure that all parameters are specified for each
selected step within a task. Press End to save and exit.

Commands: SAVE - Save parameter values
Line Commands: / - Select

Product customization library : CNK.CUST.$3090$.CNK210          More:  +

Usage Notes:
- Ensure that the DB2SORT subsystem parameter is set to ENABLE.
- For the best resource utilization, set the UTSORTAL subsystem parameter
  to YES.

Required parameters
SCNKLINK library . . . . CNK210.SCNKLINK
SCNKLPA library . . . . CNK210.SCNKLPA

/ Customize DB2 Sort

/ Change DB2 Sort installation defaults
SCNKUSER library . . . CNK210.SCNKUSER
Sort work DASD pool . . . . . SYSDA
Sort work DASD type . . . . . 3390 (3380, 3390, 9345)
Number of allocation retry attempts . . 5
Minutes between retry attempts . . . . 3
Sort work storage class . . . . .
Maximum number of sort work files . . . 32
TUNEMSG . . . . . ON (ON, OFF)

```

Figure 9. The Product Parameters panel

2. Select any required tasks and steps, and specify values for any parameters. After you select a task or step with a slash (/), put the cursor in the selected field and press Enter. If tasks, steps, and parameters are required, they are preselected with a slash (/). Otherwise, they are not preselected.

All of the required parameters have default values, which you can either accept or change.

Tips:

- In the command line, specify the KEYS command, and map EXPAND to one of the function keys.
 - For a detailed description of all input fields, put the cursor in the field, and press F1 or the key that is mapped to Help.
 - The following elements apply to specific fields:
 - **Add...** is displayed when parameters can have multiple values but currently have only one value. To specify multiple values in these fields, place the cursor on **Add...**, and press Enter. Use the displayed panel to add or delete additional values.
 - **List...** is displayed when the complete list of valid values for the fields is too long to be displayed on the panel. To see the complete list of values, place the cursor on **List...**, and press F1 or the key that is mapped to Help.
 - **More...** is displayed when input fields contains multiple values. To see all of the values in the field, place the cursor on **More...**, and press Enter.
3. Optional: Select other tasks and steps with a slash (/) and press Enter to activate the input fields. Either accept or change the default values that are displayed.
 4. Press End to save your changes and exit, or issue the SAVE command to save your changes and stay on the Product Parameters panel.

Results

The Customizer Workplace panel is displayed, and the status of the product parameters is Ready to Customize.

What to do next

If the status of other parameters on the Customizer Workplace panel is Incomplete or Discovered, edit these parameters.

Related tasks:

“Defining DB2 parameters”

DB2 parameters are parameters for a DB2 entry.

Defining DB2 parameters

DB2 parameters are parameters for a DB2 entry.

About this task

If you did not run the DB2 Sort Discover EXEC, you must create and associate one or more DB2 entries before you can define the DB2 parameters. For more information, see “Creating and associating DB2 entries” on page 40.

Procedure

1. Specify E next to one or more DB2 entries in the associated list, which is in the Associated DB2 Entries and Parameter Status section on the Customizer Workplace panel, and press Enter. The DB2 Parameters panel is displayed, as shown in the following figure:

```

CCQPDDB2          DB2 Parameters: DB2 Sort          11:34:18
Command ==>>>          Scroll ==>> CSR

Enter values for all of the DB2 parameters. Press End to save and exit.

Commands: SAVE - Save parameter values

DB2 subsystem ID . . . . . : DB02
Group attach name . . . . .

General DB2 Information - common
Mode . . . . . NFM (CM, CM8, CM9, NFM)
Level number . . . . . (810, 910, 101, 111)

DB2 Libraries
Load Library . . . . . DSN.SDSNLOAD > Add...
Run Library . . . . . DSN.RUNLIB.LOAD > Add...

DB2 Utilities - common
SYSAFF for DB2 utilities . . . . .

```

Figure 10. The DB2 Parameters panel

2. Specify values for all parameters that are displayed.

Tips:

- In the command line, specify the KEYS command, and map EXPAND to one of the function keys.
- For a detailed description of all input fields, put the cursor in the field, and press F1 or the key that is mapped to Help.
- The following elements apply to specific fields:
 - **Add...** is displayed when parameters can have multiple values but currently have only one value. To specify multiple values in these fields, place the cursor on **Add...**, and press Enter. Use the displayed panel to add or delete additional values.
 - **List...** is displayed when the complete list of valid values for the fields is too long to be displayed on the panel. To see the complete list of values, place the cursor on **List...**, and press F1 or the key that is mapped to Help.
 - **More...** is displayed when input fields contains multiple values. To see all of the values in the field, place the cursor on **More...**, and press Enter.

Many parameters have default values, which you can either accept or change.

3. Press End to save your changes and exit, or issue the SAVE command to save your changes and stay on the same panel.

Results

The status of the DB2 entries that you selected on the Customizer Workplace panel is Ready to Customize.

What to do next

If the status of other parameters on the Customizer Workplace panel is Incomplete or Discovered, edit these parameters.

Related tasks:

“Defining DB2 Sort parameters” on page 42
 DB2 Sort parameters are specific to DB2 Sort.

Generating customization jobs

To generate customization jobs for DB2 Sort and any associated DB2 entries, issue the GENERATEALL command, or select one or more DB2 entries on which to customize DB2 Sort.

Procedure

Generate the customization jobs by using one of the following methods.

- If you want to generate customization jobs at the product level and for any associated DB2 entries, issue the GENERATEALL command, and press Enter.
- If you want to generate customization jobs for specific DB2 entries, select the DB2 entries by specifying the G line command against them, and press Enter. The available DB2 entries are in the associated list in the Associated DB2 Entries and Parameter Status section.

Important: Regenerating customization jobs will replace any existing jobs, including jobs that you might have manually modified after they were generated.

Results

If the status is Incomplete or Discovered for DB2 Sort parameters or DB2 parameters, Tools Customizer automatically starts an editing session for the types of parameters that are required. The session continues until the panel for each type of required parameter has been displayed.

What to do next

If an automatic editing session is started, accept the displayed parameter values or define values for the required types of parameters, select optional parameters, tasks, or steps for your environment, and save the parameter values. Otherwise, the customization jobs are generated, and you can submit them.

Tip: If the customization jobs are generated, but you are not ready to submit them, you can see them later by issuing the JOBLIST command on the Customizer Workplace panel. The JOBLIST command displays the Finish Product Customization panel, which you can use to submit the jobs.

Submitting customization jobs

Submit the customization jobs to customize DB2 Sort.

Before you begin

Ensure that the correct jobs are generated.

If you generated an IVP job, ensure that you complete the following steps before you submit the IVP job:

- Authorize the SCNKLINK and SCNKLPA libraries.
- Issue the following DB2 command to collect performance trace data: `-START TRACE(P) DEST(SMF) CLASS(30) IFCID(25) TDATA(COR,CPU)`.

About this task

The following figure shows part of the Finish Product Customization panel. The table on this panel shows the customization jobs that are generated by Tools

Customizer. They are grouped by job sequence number.

```

CCQPCST          Finish Product Customization          Row 1 to 3 of 3
Command ==>>>                                     Scroll ==>> CSR
CCQC000I The jobs were generated on the selected DB2 entries.
For a first-time customization, submit the jobs in the members in the order
in which they apply to the DB2 entries. Otherwise, submit only the necessary
jobs that were generated after changes were made. To submit jobs, browse
the members and issue the TSO SUBMIT command.

Line Commands: E - Edit  B - Browse

      Product customization library:  CNK.CUST.$3090$.CNK210

Cmd Member  SSID GrpAttch  Template Date      Description
-----
A0RLNKJ  --  --      CNKRLNKJ 2014/06/27 Customizes DB2 Sort library
A1IVP1AA DB02 --      CNKIVP1J 2014/06/27 Creates the DB2 Sort IVP1 (10/4
A2IVP2AA DB02 --      CNKIVP2J 2014/06/27 Creates the DB2 Sort IVP2 (8/25
-----
End of customized jobs -----

```

Figure 11. The Finish Product Customization panel

The member-naming conventions depend on whether the customization jobs are for DB2 entries, and LPAR, or the product.

Customization jobs for DB2 entries

The members use the following naming convention:

<job_sequence_number><job_ID><DB2_entry_ID>

where

job_sequence_number

Two alphanumeric characters, A0 - Z9, that Tools Customizer assigns to a customization job. The number for the first template in the sequence is A0, the number for the second template is A1, and so on.

job_ID Characters 4 - 7 of the template name, if the template name contains five or more characters. Otherwise, only character 4 is used. DB2 Sort assigns the template name.

DB2_entry_ID

Two alphanumeric characters, AA - 99, that Tools Customizer assigns to a DB2 entry.

For example, the XYZBNDD*DB2_entry_ID_1* and XYZBNDD*DB2_entry_ID_2* jobs are generated from the XYZBNDGR template, and the XYZ4*DB2_entry_ID_1* and XYZ4*DB2_entry_ID_2* jobs are generated from the XYZ4 template. If the jobs are generated on two DB2 entries, the following member names are listed sequentially: A0BNDGAA, A0BNDGAB, A14AA, A14AB.

Customization jobs for the product

The members use the following naming convention:

<job_sequence_number><job_ID>

where

job_sequence_number

Two alphanumeric characters, A0 - Z9, that Tools Customizer

assigns to a customization job. The number for the first template in the sequence is A0, the number for the second template is A1, and so on.

job_ID Characters 4 - 8 of the template name, if the template name contains five or more characters. Otherwise, only character 4 is used. For example, for the XYZMAKE template, the job ID is MAKE. For the XYZM template, the job ID is M. DB2 Sort assigns the template name, and it is displayed in the Template column.

For example, the XYZBNDGR job is generated from the XYZBNDGR template, and the XYZ4 job is generated from the XYZ4 template. The following member names are listed sequentially: A0BNDGR, A14.

Procedure

1. Submit the generated customization jobs by following the process that you use in your environment or by using the following method:
 - a. Specify B against a customization job or the product customization library, and press Enter. An ISPF browsing session is started.
 - b. Browse the customization job or each member in the library to ensure that the information is correct.
 - c. Run the TSO SUBMIT command.
2. Press End.

Results

DB2 Sort is customized, and the Customizer Workplace panel is displayed. The status is Customized for the DB2 entries on which DB2 Sort was customized.

What to do next

You can generate more customization jobs for other DB2 entries, view a list of customization jobs that you previously generated, or recustomize DB2 Sort.

Browsing parameters

You can browse the product parameters and the DB2 parameters in read-only mode.

Procedure

1. On the Customizer Workplace panel, specify B next to the **Product parameters** field or the DB2 entry that you want to browse, and press Enter. The panel that corresponds to your specification is displayed.
2. Press End to exit.

Copying DB2 entries

You can copy associated and not associated DB2 entries to other DB2 entries or to new DB2 entries.

About this task

Go to the step that applies to your environment:

- To copy an associated DB2 entry to another associated DB2 entry or to an entry that is not associated, go to step 1.
- To copy an associated DB2 entry to a new entry, go to step 2.

- To copy a DB2 entry that is not associated to a new entry, go to step 3.

Procedure

1. To copy an associated DB2 entry to another associated DB2 entry or to an entry that is not associated, complete the following steps:
 - a. Specify C against a DB2 entry in the associated list of DB2 entries on the Customizer Workplace panel, and press Enter. The Copy Associated DB2 Entry panel is displayed.
 - b. Select one or more DB2 entries to which information will be copied by specifying the / line command, and press Enter. The Associated column indicates whether the DB2 entry is associated.

Tip: To copy information into all of the DB2 Entries in the list, issue the SELECTALL primary command, and press Enter. The Copy DB2 Parameter Values panel is displayed.

- c. Specify an option for copying common and product-specific DB2 parameter values. Common DB2 parameter values apply to all DB2 entries for all products that you have customized by using Tools Customizer. Product-specific DB2 parameter values apply only to the product that you are currently customizing.
 - To copy the common DB2 parameter values and the product-specific DB2 parameter values, specify option 1, and press Enter.
 - To copy only the product-specified DB2 parameter values, specify option 2, and press Enter.

In some cases, the DB2 parameter values might contain the DB2 subsystem ID as an isolated qualifier in data set names. For example, in the DB01.DB01TEST.DB01.SANLLOAD, data set name, the DB01 subsystem ID is isolated in the first and third qualifiers but is not isolated in the second qualifier. When the DB2 subsystem ID is an isolated qualifier in data set names, the Change DB2 Subsystem ID in DB2 Parameter Values panel is displayed. Otherwise, the Customizer Workplace panel is displayed.

- d. If the Change DB2 Subsystem ID in DB2 Parameter Values panel is displayed, specify an option for changing the subsystem IDs. Otherwise, skip this step.
 - To change the subsystem ID in isolated qualifiers in data set names, specify option 1, and press Enter.
 - To use the same subsystem ID in all values, specify option 2, and press Enter.

The Customizer Workplace panel is displayed with the copied associated entry in the list.

2. To copy an associated DB2 entry to a new entry, complete the following steps:
 - a. Specify C against a DB2 entry in the associated list of DB2 entries on the Customizer Workplace panel, and press Enter. The Copy Associated DB2 Entry panel is displayed.
 - b. Issue the CREATE command. The Create DB2 Entries panel is displayed.
 - c. Specify the SSID, the group attach name, or both in the appropriate columns for each new DB2 entry, and press Enter.

Tip: To add rows for additional entries, specify the *Imm* line command, where *nn* is the number of entries to be created, and press Enter. The Copy Associated DB2 Entry panel is displayed with the new entries in the list. The new entries are preselected.

- d. Press Enter to complete the copy process. The Customizer Workplace panel is displayed with the copied entries in the list.
3. To copy a DB2 entry that is not associated to a new entry, complete the following steps:
 - a. Issue the ASSOCIATE command on the Customizer Workplace panel. The Associate DB2 Entry for Product panel is displayed.
 - b. Select one or more DB2 entries by specifying the / line command, and press Enter. The Copy a DB2 Entry panel is displayed.
 - c. Specify the SSID, the group attach name, or both in the appropriate columns for the new DB2 entry, and press Enter. The Associate DB2 Entry for product panel is displayed with the copied entry in the list.
 - d. If you want to associate the copied entry, specify A against it, and press Enter. The Customizer Workplace panel is displayed with the copied entries in the list.

What to do next

Edit any of the parameters or generate the jobs.

Related concepts:

“Tools Customizer terminology” on page 18

Tools Customizer uses several unique terms that you should be familiar with before you begin to use Tools Customizer.

Removing DB2 entries

You can remove DB2 entries from the associated list.

About this task

When you remove DB2 entries from the associated list, any customization jobs for the entries are removed from the list of jobs on the Finish Product Customization panel, and they are deleted.

Procedure

On the Customizer Workplace panel, specify R next to one or more DB2 entries that you want to remove, and press Enter. The selected DB2 entries are removed from the associated list and added to the master list on the Associate DB2 Entry for Product panel, and the customization jobs are deleted.

Related concepts:

“Tools Customizer terminology” on page 18

Tools Customizer uses several unique terms that you should be familiar with before you begin to use Tools Customizer.

Deleting DB2 entries

You can delete DB2 entries from the master list.

About this task

When you delete DB2 entries from the master list, any associations and all customization jobs for products that are customized on the entries will be deleted.

Procedure

1. On the Customizer Workplace panel, issue the ASSOCIATE command. The Associate DB2 Entry for Product panel is displayed.
2. Specify D next to one or more DB2 entries that you want to delete, and press Enter. If the entry is associated with any products, the Delete Associated DB2 Entry panel for the first DB2 entry that you selected is displayed. Otherwise, the Delete DB2 Entry panel is displayed.
3. To delete the DB2 entries, press Enter. If the DB2 entries are associated with any products in the table on the Delete Associated DB2 Entry panel, any associations and all customization jobs for the products that are customized on it are deleted. Otherwise, only the DB2 entries are deleted. If you selected multiple DB2 entries to delete, the next DB2 entry that you selected is displayed on either the Delete Associated DB2 Entry panel or the Delete DB2 Entry panel. Otherwise, the Associate DB2 Entry for Product panel is displayed.

What to do next

If you selected multiple DB2 entries to delete, repeat step 3 until all selected entries are deleted. Then, continue the customization process.

Displaying customization jobs

You can view a list of the members that contain the customization jobs before or after you submit the jobs.

About this task

The customization jobs that you generate for one DB2 entry are also displayed when you customize DB2 Sort for another DB2 entry later.

Procedure

On the Customizer Workplace panel, issue the JOBLIST command. The Finish Product Customization panel is displayed. This panel shows the list of jobs that you have previously generated. They are grouped by job sequence number. Use this panel to browse or edit the generated jobs before you submit them.

Maintaining customization jobs

Instead of deleting customization jobs outside of Tools Customizer, you can maintain the correct jobs for DB2 Sort by completing the steps for recustomization.

About this task

You cannot delete or rename customization jobs from the customization library by starting an ISPF browse or edit session from the Finish Product Customization panel. If you try to delete customization jobs by using this method, the CCQC034S message is issued. If you try to rename customization jobs, the CCQC035S message is issued.

If you delete or rename customization jobs from the customization library by using ISPF outside of Tools Customizer, Tools Customizer will not recognize that the jobs were deleted, and the Finish Product Customization panel will still display them. If you browse or edit jobs that were deleted from the library outside of Tools Customizer, the CCQC027S message is issued.

Procedure

To maintain the correct customization jobs in the customization library, complete the steps for recustomization.

Using Tools Customizer in a multiple-LPAR environment

Currently, Tools Customizer supports only the local LPAR; however, you can propagate customizations to additional LPARs by using either of two different methods.

About this task

In a multiple-LPAR environment, Tools Customizer identifies the LPAR to which you are logged on. Tools Customizer uses this LPAR name for several different parameter settings, one of which is the data store. When you use the data store during the customization of DB2 Sort that is on a different LPAR, Tools Customizer issues message CCQD586S, which indicates that the product has already been customized based on values from the data store on the first LPAR. This message is issued to prevent the data store from becoming corrupted.

This behavior occurs in the following conditions:

- Tools Customizer is installed on a DASD device that is shared by multiple LPARs.
- After a product is customized by using Tools Customizer, the data store is copied to another LPAR.

Procedure

To customize products running against a DB2 subsystem on an LPAR where Tools Customizer is not installed, consider using one of the following methods:

Install one instance of Tools Customizer on one LPAR

If you intend to reuse the customization values for all the instances of your products on all LPARs, use this method.

1. Associate all the DB2 entries in this one instance of Tools Customizer. The LPARs on which the DB2 subsystems reside do not matter.
2. Generate the customization jobs for each DB2 entry.
3. Copy the generated customization jobs to the LPAR to run against the specific DB2 entries. Some LPAR-specific edits might be required. You can make these edits in the customized jobs that you copied. Note that this situation is one of the few situations where you might need to make manual changes to the jobs that are customized by Tools Customizer.

Install one instance of Tools Customizer on each LPAR

If you do not want to reuse previous customization values and you want to start new customizations, use this method.

Important: This method will likely not be the preferred approach for most organizations because most organizations tend to use similar or identical customization values for each product instance on all LPARs.

Accessing DB2 Sort libraries

After activating and customizing DB2 Sort, you must create access to the following DB2 Sort libraries: SCNKLINK and SCNKLPA. If access to these two libraries was established before you began this customization, then ignore the procedure that follows and refresh the library by using the following MVS™ system command:
MODIFY LLA,REFRESH

Procedure

Decide how you want to access the DB2 Sort libraries, SCNKLINK and SCNKLPA. To make these decisions, consider the following information about the attributes of the DB2 Sort libraries:

- The SCNKLINK and SCNKLPA libraries can be in the LNKST.
- The SCNKLPA library can be in the LPA, but it is not required to be in the LPA.
- If the SCNKLPA library is not in the LPA, it must be APF-authorized.
- The SCNKLINK library must be APF-authorized.

What to do next

You can submit the Installation Verification Program job that you generated with Tools Customizer to test your product configuration before moving to a production environment. After the IVP job runs to completion, refer to the procedure to generate an IVP report in “Generating an IVP report” on page 54.

After the product has been successfully installed and configured and the load library is APF-authorized, you may begin using DB2 Sort. Refer to the sample job in “Running with DB2 Sort” on page 55.

Chapter 4. Verifying installation and customization

Before using DB2 Sort in a production environment, you can test the installation and customization in a test environment using the Installation Verification Program (IVP).

Topics:

- “Installation Verification Program: Overview”
- “Using Installation Verification Program” on page 54
- “Generating an IVP report” on page 54

Installation Verification Program: Overview

Installation Verification Program (IVP) is a program you can run in a test environment to verify your DB2 Sort installation prior to moving to a production environment.

DB2 Sort IVP is a self-contained evaluation and testing methodology you can use to determine that DB2 Sort is installed correctly and performing properly.

When you run IVP, it performs the following tasks:

- Creates a storage group and database for the IVP objects and then drops them at the end of the job
- Creates tables with different sizes and row lengths, and multiple indexes and index lengths
- Ensures data is not in clustering sequence
- Executes the LOAD, REBUILD INDEX, and REORG utilities with and without DB2 Sort

When you are customizing DB2 Sort with Tools Customizer, you can choose the cases for IVP. There are four table sizes: 8 billion bytes, 10 billion bytes, 25 billion bytes and 40 billion bytes. Base your decision on the amount of disk space and time you want the job to use.

To ensure that the IVP results are meaningful, you must run IVP in an appropriate test environment that represents your production environment as much as possible.

Use the following general requirements when setting up your test environment:

- **CPU:** 2 general purpose processors minimum, preferably 4 or more; 1 or more zIIP processors
- **DASD:** Minimum of 6 volumes (PAV-capable)
- **Channels:** Minimum of 6 (quiet)
- **Memory:** Minimum of 10 GB
- **LPAR:** Stand-alone or quiet system
- **Traces:** Turn off all DB2 Utility traces that impact performance times

Using Installation Verification Program

Use the Installation Verification Program (IVP) in a test environment to verify that your DB2 Sort installation and any customizations you made are functioning correctly before moving to a production environment.

Before you begin

Make sure that DB2 Sort has been installed and you have an appropriate test environment. For details about setting up the test environment, see “Installation Verification Program: Overview” on page 53.

Since information about the IVP job will be recorded in SMF records and performance trace data, you must ensure that your system is set up to generate the performance trace data, using the following requirement:

- Issue the following DB2 command to collect performance trace data:
-START TRACE(P) DEST(SMF) CLASS(30) IFCID(25) TDATA(COR,CPU)

Note: Issue a DISPLAY TRACE(*) first. If another active trace with CLASS(30) has started, you will need to issue a START TRACE(P) CLASS(31) or (32). If all three classes are in use, add IFCID(25) using a MODIFY TRACE command. Alternatively, you can issue the STOP TRACE(P) command and reissue the START TRACE(P), combining the different IFCIDs into the single START TRACE(P) command.

If you have generated the IVP job with Tools Customizer, submit the job after you have set up your environment to generate the performance trace data and after your DB2 Sort libraries have been made accessible. For instructions on making your libraries accessible, see “Accessing DB2 Sort libraries” on page 52.

Generating an IVP report

After you run one of the IVP jobs, you can generate a report to display the results.

About this task

The report program uses SMF type 30 subtype 4 and DB2 SMF 102 (IFCID 25) records as input.

To display the spreadsheet located in the SCNKBENU library, a PC/workstation with a program that can edit spreadsheets from Microsoft Office Excel 2003 (XLS format) or later is required.

Procedure

1. Extract the SMF type 30 subtype 4 records and the SMF 102 records from your SMF data sets and put them into a single data set. You can use the IFASMFDP program to create the data set.
2. If DB2 SMF records are in compressed format, you must decompress the SMF records prior to using the reporter job with the SMF decompression program DSNTSMFD.

Note: For more information about the SMF decompression program DSNTSMFD, refer to:

- *DB2 11 for z/OS Utility Guide and Reference (SC19-4067)*
- *DB2 10 for z/OS Utility Guide and Reference (SC19-2984)*

- For DB2 10 the SMF decompression program is provided with PTF UK64597.
3. Open the CNKIVRPT member which is located in the SCNKXSAMP library.
 4. Make the editing changes according to the instructions that are included in the JCL and submit the job. This job creates an output file named *userid*.CNKDATA.TXT.
 5. Transfer your *userid*.CNKDATA.TXT file to a PC/workstation in ASCII mode, renaming it to *c:\temp\cnkdata.txt*.

Note: If your transfer facility offers a carriage return/ line feed option (CRLF), it should be selected.

6. Transfer the member named CNKIVXLS which is located in the SCNKBENU library to a PC/workstation in binary mode, renaming it to *cnkivxls.xls*.
7. On the PC/workstation, open the *cnkivxls.xls* spreadsheet. Enable content in response to the security warning. Select the Worksheet tab and follow the directions within the worksheet.

Results

The IVP report program produces the following two reports:

- Summary
- Details

After the worksheet has been filled in, choose the Summary and Details tabs to view the reports.

What to do next

After the product has been successfully installed and configured and the load library is APF-authorized, you may begin running DB2 utilities and tools with DB2 Sort.

Running with DB2 Sort

For testing purposes, specifying the DB2 Sort libraries in the STEPLIB may be the simplest and most straightforward way to get started. The utility and tool jobs need to be updated to include the DB2 Sort libraries in the STEPLIB.

Example

Here is a sample job with the DB2 Sort load libraries specified in the STEPLIB of the job:

```
//SRTREORG JOB ('DB2/APPL'),'REORG00',REGION=0M
...
//STEP1 EXEC PGM=DSNUTILB,PARM=(DBM0,'SRZTAUM3'),REGION=0M
...
//STEPLIB DD DSN=DB2M.SDSNEXIT,DISP=SHR
//          DD DSN=DB2M.SDSNLOAD,DISP=SHR
//          DD DSN=<hlq>.DB2SORT.SCNKLINK,DISP=SHR
//          DD DSN=<hlq>.DB2SORT.SCNKLPA,DISP=SHR
...
//SYSPRINT DD SYSOUT=*
//UTPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
...
```

Chapter 5. Troubleshooting DB2 Sort

Use these topics to diagnose and correct problems that you experience with DB2 Sort.

Topics:

- “DB2 Sort messages”
- “Gathering diagnostic information” on page 70

DB2 Sort messages

The information in the DB2 Sort messages can help you diagnose and resolve DB2 Sort problems.

Message format

DB2 Sort messages adhere to the following format:

CNKnnnx

Where:

CNK Indicates that the message was issued by DB2 Sort

nnn Indicates the message identification number

x Indicates the severity of the message:

E Indicates that an error occurred, which might or might not require operator intervention

I Indicates that the message is informational only

Each message also includes the following information:

Explanation:

The Explanation section explains what the message text means, why it occurred, and what its variables represent.

System action:

The System action section explains what the system will do in response to the event that triggered this message.

User response:

The User response section describes whether a response is necessary, what the appropriate response is, and how the response will affect the system or program.

DB2 Sort messages

The following DB2 Sort messages are listed in numerical order, from lowest to highest.

CNK036I G=ggg, B=bbb, SEGMENT
LENGTH=sss, BIAS=zz

Explanation: This message displays the following tuning information: ggg is the number of records that

can be contained in the working virtual storage area that is used by DB2 Sort V2.1. For variable-length records, this number is the number of segments. The bbb value indicates the physical blocking that is used for intermediate storage. For fixed-length records, this

CNK046E • CNK071E

number represents the blocking factor. For variable-length records, it represents the block size. This value is excluded for incore or turnaround sorts. The sss value indicates the segment length that is used in the working storage that DB2 Sort V2.1 uses during Phase 1. This value is included in the message for variable-length records when the execution is not an incore or turnaround sort. The zz value indicates the degree of prior sequencing in the input data. This number ranges from 00 to 99. A low value indicates random input, and a high value indicates highly sequenced input. The BIAS value is not included in the message for an incore or turnaround sort in which it is 100 by definition.

System action: None.

User response: No action is required.

CNK046E SORT CAPACITY EXCEEDED

Explanation: All available intermediate storage (sort work space) is exhausted, including any secondary allocation. Sort processing cannot continue.

System action: Processing stops.

User response: If fewer than 255 data sets have been allocated, specify a sufficient number in the **DYNALLOC** runtime parameter in the DB2SORTI data set, and then resubmit the job. If 255 data sets have been allocated, ask the system programmer who is responsible for installing DB2 Sort to add disk space to the class of temporary data sets that is used to allocate sort work data sets.

CNK052I END DB2 SORT FOR Z/OS – jobname, stepname, procstepname, DIAG=hhhh,hhhh,...

Explanation: DB2 Sort completed successfully. The hexadecimal information that follows the DIAG keyword is likely to change from execution to execution. It is internal diagnostic information intended for use by IBM Software Support.

System action: None.

User response: No action is required.

CNK054I RECORDS IN x, OUT y

Explanation: This message provides information about record quantities. x represents the number of records that were read from the input data set(s). y represents the number of records that were written to the output file.

System action: None.

User response: No action is required.

CNK055I INSERT x, DELETE y

Explanation: This message provides information about insert and delete activity. x represents the number of records that were inserted by the input exit routine. y represents the number of records that were deleted by the output exit routine.

System action: None.

User response: No action is required.

CNK061E I/O ERROR jobname, stepname, unit address, device type, DDname, operation attempted, error description, last seek address or block count, access method.

Explanation: An I/O error occurred on the device whose address is given.

System action: Processing stops.

User response: Resubmit the job because I/O errors are often transient. If the job continues to fail, contact the system programmer who is responsible for installing DB2 Sort.

CNK063E xxxxxx OPEN ERROR

Explanation: The data set shown cannot be successfully opened.

System action: Processing stops.

User response: Ask the system programmer who is responsible for installing DB2 Sort to contact IBM Software Support.

CNK066E APPROXIMATE RECORD COUNT x

Explanation: DB2 Sort stopped because an insufficient amount of intermediate storage was available, as indicated in the accompanying CNK046E message. x indicates the approximate number of records that were processed by DB2 Sort up to this point.

System action: Processing stops.

User response: See the user response for "CNK046E."

CNK071E MAXIMUM NUMBER OF RECORDS EXCEEDED

Explanation: The number of records in this application is greater than the default internal limit.

System action: Processing stops.

User response: Ask the system programmer who is responsible for installing DB2 Sort to contact IBM Software Support.

CNK072I {EQUALS,NOEQUALS}

Explanation: This message provides tuning information. If EQUALS is in effect, the input order of records with equal keys is preserved. If NOEQUALS is in effect, records with equal keys might not be in the same order as they were presented to the sort or merge process.

System action: None.

User response: No action is required.

CNK073I ddname: dsname [(FIRST of n)]

Explanation: This message displays the SORTIN DD data set name. For concatenated DDs, only the first data set name will be displayed and the number of concatenations will be displayed in "FIRST of n".

System action: None.

User response: No action is required.

CNK074I ddname: dsname

Explanation: This message displays the SORTOUT DD data set name.

System action: None.

User response: No action is required.

CNK106E ddname INVALID DEVICE TYPE

Explanation: The input or output file resides on an invalid device type. ddname is SORTIN or SORTOUT. Valid device types include the IBM 3380, 3390, and 9345 direct access devices and their equivalents as well as the IBM 3420, 3480, 3490, and 3590 series tape devices and their equivalents.

System action: Processing stops.

User response: Copy the data set to a valid device and resubmit the job.

CNK108I ddname: RECFM= ;LRECL= {;BLKSIZE=,CISIZE=} [;CINV ACCESS]

Explanation: This message displays the DCB characteristics used by DB2 Sort to process the input file. For a non-VSAM data set that is concatenated, the DCB characteristics are for the first of the concatenated data sets, except for BLKSIZE, which is the largest block size of all data sets in the concatenation examined at sort initialization time. For a VSAM data set, the CISIZE is provided; if control interval access was used, the CINV ACCESS portion of the message will be displayed.

System action: None.

User response: No action is required.

CNK109I MERGE INPUT: TYPE={E,V};LRECL=

Explanation: This message lists the DCB (Data Control Block) characteristics that are used by DB2 Sort to process the input files for a merge.

System action: None.

User response: No action is required.

CNK110I ddname RECFM= ;LRECL= {;BLKSIZE=,CISIZE=} [;CINV ACCESS]

Explanation: This message displays the DCB characteristics used by DB2 Sort to process the output file. For a VSAM data set, the CISIZE is provided; if control interval access was used, the CINV ACCESS portion of the message will be displayed.

System action: None.

User response: No action is required.

CNK122E INVALID INTERMEDIATE STORAGE DEVICE

Explanation: A device that is not valid was assigned as intermediate storage. Valid devices include the 3380, 3390, and 9345 mass storage systems, and equivalent units.

System action: Processing stops.

User response: If you specified a **DYNALLOC** parameter in a DB2SORTI data set, ensure that a valid device was specified. Specify a valid device and resubmit the job. If you did not specify a **DYNALLOC** runtime parameter, ask the system programmer who is responsible for installing DB2 Sort to change the class of temporary data sets that are used to allocate sort work data sets so that only valid devices are available.

CNK124I [ESTIMATED] PREALLOCATED/USED SORTWORK SPACE USAGE FACTOR {=,<,>}nn.nn

Explanation: This message provides information about the amount of sort work space that was allocated and the amount of sort work space that was actually used. nn.nn represents the quotient that is obtained by dividing the number of tracks that are assigned within pre-allocated sort works by the number of tracks that are actually used by DB2 Sort. The word ESTIMATED is included when the derivation of this factor is inexact, for example, when all sort work data sets are not opened, or when data space or memory objects are used to contain part or all of the sort work data.

System action: None.

User response: No action is required.

CNK130E I/O ERROR ON SYSIN

Explanation: An I/O error occurred on the options override input DD statement.

System action: Processing stops.

User response: Check the DD statement to ensure that the data set is properly allocated, and then resubmit the job.

CNK131I PARM FIELD ERROR – xxxxxxxx

Explanation: An invalid PARM was in the PARM field string that was passed to DB2 Sort.

System action: PARM processing stops, the remainder of this PARM string is ignored, and the application continues to run. The entire PARM string is written to the messages data set, and an asterisk is displayed beneath the invalid PARM.

User response: DB2 Sort might appear to have completed successfully, but because the invalid PARM and any PARMs after it were ignored, the application might not have produced the anticipated results. If the anticipated results were not achieved, correct the PARM and resubmit the job.

CNK133E Exx USER EXIT REQUESTED TERMINATION

Explanation: The DB2 utility encountered a critical problem while processing, so passed a return code directing DB2 Sort to stop processing.

System action: Processing stops.

User response: Review the DB2 utility-generated messages prior to the CNK133E which indicate a processing problem. Resolve this problem and resubmit the job.

CNK135I TASK CALL/E35 TERMINATED PREMATURELY

Explanation: An E35 exit routine passed a return code of 8, which stopped DB2 Sort before it was able to pass all of the records to the sort.

System action: None.

User response: Ask the system programmer who is responsible for installing DB2 Sort to contact IBM Software Support.

CNK138E ddname BLKSIZE NOT EVENLY DIVISIBLE BY LRECL

Explanation: The length of a block from the input or output file was not a multiple of the LRECL value, or the JCL or data set attributes are incorrect. ddname is SORTIN or SORTOUT.

System action: Processing stops.

User response: Correct the data set attribute specifications and resubmit the job.

CNK141E ddname RECFM IS U

Explanation: DB2 Sort does not support an undefined record format. ddname is SORTIN or SORTOUT.

System action: Processing stops.

User response: Specify a data set that does not have an undefined record format and resubmit the job.

CNK142E MIXED FIXED AND VARIABLE SORTIN TYPES NOT SUPPORTED

Explanation: DB2 Sortt permits only one record format type (fixed or variable) for input files per sort.

System action: Processing stops.

User response: Correct the data set attribute specifications and resubmit the job.

CNK144I UNEXPECTED VIRTUAL STORAGE FRAGMENTATION

Explanation: The amount of virtual storage that was calculated by DB2 Sort for Phase 2 or 3 was not available in a contiguous block. Additional virtual storage was obtained to satisfy the requirement.

System action: None.

User response: No action is required.

CNK146I nnn BYTES OF EMERGENCY SPACE

Explanation: nnn indicates the amount of virtual storage that has been set aside by DB2 Sort for use by other programs (for example, the caller, system SVCs, or the tape management system).

System action: None.

User response: No action is required.

CNK149I FRAGMENTED VIRTUAL STORAGE IN SORT PHASE

Explanation: The amount of virtual storage that is specified for use by DB2 Sort was not available in a contiguous block for Phase 1. DB2 Sort obtained its virtual storage in fragments and continued execution.

Note: The calling program or user exit routine used virtual storage in a way that caused fragmentation, which might result in ABEND S80A or S804 when you run subsequent DB2 Sort jobs.

System action: None.

User response: No action is required.

**CNK151I SECONDARY EXTENTS OBTAINED
 xxx**

Explanation: xxx indicates the number of secondary extents that were obtained for sort work data sets.

System action: None.

User response: No action is required.

**CNK152I REQUESTED VIRTUAL STORAGE
 NOT AVAILABLE, nnn BYTES USED**

Explanation: The amount of virtual storage requested by the caller was not available when DB2 Sort received control. nnn indicates the number of available bytes that were used by DB2 Sort.

System action: None.

User response: No action is required.

**CNK153E SEVERELY UNDERESTIMATED
 FILESIZE**

Explanation: The file size estimate that was provided to the utility execution was substantially lower than the actual file size. DB2 Sort used this estimate to determine the amount of resources to allocate. Because the real file size was much larger than the estimate, the amount of resources allocated was insufficient for the sort process to complete.

System action: Processing stops.

User response: Correct the file size estimate and resubmit the job.

**CNK157E SPANNED RECORD LENGTH
 LARGER THAN LRECL**

Explanation: A record from an input data set with RECFM=VBS contains a record longer than the maximum record length specified by the LRECL in the DCB.

System action: Processing stops.

User response: Correct the data set attribute specifications and resubmit the job.

**CNK159E RECORD LENGTH 0, SORTIN
 RECORD x**

Explanation: An invalid variable-length record (length code <4 in its Record Descriptor Word) has been found in the input file. x is the number of the invalid record.

System action: Processing stops.

User response: Correct the data set and resubmit the job.

**CNK160E RECORD LENGTH GREATER THAN
 LRECL**

Explanation: A variable-length record read from the input file is longer than the maximum record length specified by the LRECL in the DCB.

System action: Processing stops.

User response: Correct the data set and resubmit the job.

CNK161I ALTERNATE PARM USED

Explanation: The alternate PARM option data set was used, and DB2 Sort received the specified parameters.

System action: None.

User response: No action is required.

**CNK162I ppp PREALLOCATED SORTWORK
 TRACKS, ddd DYNAMICALLY
 ALLOCATED sss ACQUIRED IN xxx
 SECONDARY EXTENTS, rrr
 RELEASED, TOTAL OF uuu TRACKS
 USED**

Explanation: This message provides information about the number of tracks that were used by DB2 Sort. ppp is the number of tracks that were found available in sort work data sets that were allocated prior to DB2 Sort gaining control. (These data sets might have been dynamically allocated by the caller.) ddd is the number of tracks that were dynamically allocated as primary space by DB2 Sort. sss is the number of tracks that were acquired as secondary space on both pre-allocated data sets and data sets that were dynamically allocated by DB2 Sort. xxx is the total number of secondary extents that were acquired. rrr is the total number of unneeded tracks that were released from both pre-allocated data sets and data sets that were dynamically allocated by DB2 Sort. uuu is the total number of tracks that were actually used in sorting. The following notes apply to the information in this message: ppp might not represent all of the pre-allocated tracks that were available because not all preallocated sort work data sets might be opened by DB2 Sort. uuu might be less than the sum of ppp, ddd and sss because uuu represents the space that was actually used and not the space that was available.

System action: None.

User response: No action is required.

**CNK164I www BYTES OF VIRTUAL STORAGE
 AVAILABLE, xxx BYTES REQUESTED,
 yyy BYTES RESERVE REQUESTED, zzz
 BYTES USED**

Explanation: This message reflects the total amount of virtual storage below and above the 16-MB line that was available to DB2 Sort and used by DB2 Sort. www

CNK167E • CNK180E

indicates the amount of virtual storage that was available when DB2 Sort received control. xxx indicates the amount of virtual storage that was requested for use by DB2 Sort. yyy indicates the amount of virtual storage that the caller requested DB2 Sort to reserve below the 16-MB line. zzz indicates the amount of virtual storage that was used by DB2 Sort.

System action: None.

User response: No action is required.

CNK167E RECORD LENGTH GREATER THAN OUTPUT LRECL

Explanation: DB2 Sort has encountered a variable-length record longer than the LRECL in the output file's DCB.

System action: Processing stops.

User response: Correct the data set and resubmit the job.

CNK170E CONCATENATED DATA SET, BLKSIZE NOT DIVISIBLE BY LRECL

Explanation: One of the files concatenated to a fixed-length input data set has a BLKSIZE that is not evenly divisible by the original LRECL.

System action: Processing stops.

User response: Correct the data set attribute specifications and resubmit the job.

CNK171E CONCATENATED DATA SET, LRECLS NOT EQUAL OR RECFMS DIFFERENT

Explanation: One of the files concatenated to a fixed-length input data set has an LRECL not equal to the original LRECL; or one of the files concatenated to a variable-length data set has an LRECL greater than the original LRECL; or one of the files concatenated to a fixed or variable-length data set has a RECFM not equal to the original RECFM.

System action: Processing stops.

User response: Correct the data set attribute specifications and resubmit the job.

CNK172E CONCATENATED DATA SET, BLKSIZE GREATER THAN ORIGINAL BLKSIZE

Explanation: One of the files concatenated to an input data set has a BLKSIZE greater than the original BLKSIZE.

System action: Processing stops.

User response: Correct the data set attribute specifications and resubmit the job.

CNK173E BDW INVALID

Explanation: The Block Descriptor Word of a block in the input data set contains a value less than 8; or the Block Descriptor Word contains a value greater than the number of bytes actually read.

System action: Processing stops.

User response: Correct the data set and resubmit the job.

CNK174E RDW INVALID, OVERFLOWS BUFFER

Explanation: The Record Descriptor Word of a record in the input data set is too large. (According to the RDW, the record extends beyond the buffer.)

System action: Processing stops.

User response: Correct the data set and resubmit the job.

CNK177I TURNAROUND SORT PERFORMED

Explanation: DB2 Sort was able to sort the input file without using intermediate storage (sort work data sets). All input data was contained in virtual storage.

System action: None.

User response: No action is required.

CNK178E ddname [nnnnn] MEMBER NOT FOUND

Explanation: An input DD statement specified a member of a partitioned data set that could not be found. nnnnn represents the concatenation number of the data set that has the member-not-found condition if the data set is part of a concatenation.

System action: Processing stops.

User response: Correct the data set specification and resubmit the job.

CNK179E ddname INVALID DCB PARAMETERS

Explanation: DB2 Sort was unable to derive RECFM, LRECL, and BLKSIZE parameters from the JCL, the DSCB on the disk or the tape label.

System action: Processing stops.

User response: Supply the data set attributes and resubmit the job.

CNK180E ddname MEMBER NOT SPECIFIED

Explanation: An input or output DD statement defines a partitioned data set, but a member name has not been specified. ddname is SORTIN or SORTOUT.

System action: Processing stops.

User response: Specify a member name on the indicated DD statement or change the partitioned data set to a sequential data set, and resubmit the job.

CNK182E INVALID RDW ddname BLOCK x

Explanation: An invalid spanned record indicator was detected in an input file whose RECFM=VBS. x is the number of the block in the file.

System action: Processing stops.

User response: Correct the data set and resubmit the job.

**CNK185I ddname DCB BLKSIZE GREATER
 THAN ACTUAL BLKSIZE, I/O
 INEFFICIENT**

Explanation: The I/O rate is reduced to an inefficient level because the block size specified for the input data set is larger than the actual block size, causing excessive error correction.

System action: None.

User response: No action is required. However, to improve I/O efficiency, correct the block size specification before resubmitting the job.

**CNK187E ddname CINV SIZE LESS THAN
 RECORD LENGTH BUT SPANNING
 NOT SPECIFIED**

Explanation: The record length is greater than the control interval size specified in the definition of the indicated VSAM data set, but the data set definition did not also include a specification for spanned records.

System action: Processing stops.

User response: Correct the data set and resubmit the job.

**CNK188E ddname IS DIRECT ACCESS/DSCB
 NOT FOUND/OBTAIN FAILED**

Explanation: DB2 Sortt was unable to successfully issue an OBTAIN for the specified direct access data set and was therefore unable to determine the DCB characteristics for the file. The OBTAIN failed either because the volume parameter was incorrectly specified for the input file indicated or because the data set was deleted from the volume. (NOTE: The data set may still be in the master catalog even though the data set is no longer on the volume.)

System action: Processing stops.

User response: Correct the data set specification and resubmit the job.

CNK189E ddname DCB RECFM REQUIRED

Explanation: The RECFM was not specified on the indicated input DD statement, nor was it available in the DSCB on disk or the tape label.

System action: Processing stops.

User response: Specify the RECFM on the DD statement and resubmit the job.

CNK191E ddname BLKSIZE/LRECL INVALID

Explanation: This message indicates that the BLKSIZE or LRECL of the input data set is invalid. The accompanying CNK108I message displays the DCB characteristics of the data set. BLKSIZE and LRECL must be equal if RECFM=F. BLKSIZE must be evenly divisible by LRECL if RECFM=FB. BLKSIZE must be greater than or equal to LRECL + 4 if RECFM=V.

System action: Processing stops.

User response: Correct the DCB specifications and resubmit the job.

CNK192E ddname DCB LRECL MISSING

Explanation: The LRECL was not specified on the indicated input DD statement, in the DSCB on the disk, nor in the tape label.

System action: Processing stops.

User response: Specify the LRECL on the DD statement and resubmit the job.

**CNK193E ddname DCB LRECL AND BLKSIZE
 MISSING**

Explanation: The BLKSIZE or LRECL must be specified on the indicated input DD statement, in the DSCB on the disk, or in the tape label.

System action: Processing stops.

User response: Specify the BLKSIZE or the LRECL on the DD statement and resubmit the job.

**CNK194E SORTOUT DCB REQUIRED/TAPE NOT
 SL**

Explanation: DISP=OLD was specified on the SORTOUT DD statement, the tape label was not specified as SL in the LABEL parameter, and required DCB information (LRECL, RECFM, BLKSIZE) was not specified.

System action: Processing stops.

User response: Specify the DCB attributes on the DD statement and resubmit the job.

CNK195E ddname DCB REQUIRED/VSAM INPUT

Explanation: The indicated output file requires additional DCB information (RECFM, LRECL, or BLKSIZE) on its DD statement.

System action: Processing stops.

User response: Specify the DCB attributes on the DD statement and resubmit the job.

CNK196E ddname RECFM=VB, LRECL GREATER THAN BLKSIZE

Explanation: RECFM=VB requires the BLKSIZE of the input ddname to be greater than or equal to LRECL + 4.

System action: Processing stops.

User response: Correct the DCB attributes and resubmit the job.

CNK197E ddname RECFM=F/FB, LRECL/BLKSIZE INVALID

Explanation: BLKSIZE and LRECL were not equal on the indicated DD statement for RECFM=F, or BLKSIZE was not a multiple of LRECL for RECFM=FB.

System action: Processing stops.

User response: Correct the DCB attributes and resubmit the job.

CNK198E ddname VARIABLE LRECL LESS THAN OR EQUAL TO 4

Explanation: The LRECL specification on the indicated input or output DD statement did not allow 4 bytes for the RDW plus 1 byte for data.

System action: Processing stops.

User response: Correct the LRECL specification and resubmit the job.

CNK199E ddname RECORD TYPE=V, BLKSIZE LESS THAN OR EQUAL TO 8

Explanation: The BLKSIZE specified for the indicated input or output DD statement did not allow 4 bytes for the BDW, 4 bytes for the RDW plus 1 byte of data.

System action: Processing stops.

User response: Correct the BLKSIZE specification and resubmit the job.

CNK200E ddname RECFM=V/VB LRECL/BLKSIZE INVALID

Explanation: RECFM=V or VB requires the BLKSIZE to be greater than or equal to LRECL + 4.

System action: Processing stops.

User response: Correct the RECFM or BLKSIZE specification and resubmit the job.

CNK201E ddname is DIRECT ACCESS/DSCB NOT FOUND/OBTAIN FAILED

Explanation: DB2 Sort was unable to successfully issue an OBTAIN for the specified direct access data set, and was therefore unable to determine the DCB characteristics of the indicated output file. The OBTAIN failed either because the volume parameter was incorrectly specified for the indicated output file, or because the data set was deleted from the volume. (NOTE: The data set name may still be in the master catalog even though the data set is no longer on the volume.)

System action: Processing stops.

User response: Correct the data set specification and resubmit the job.

CNK202E ddname RECFM INCOMPATIBLE

Explanation: The record format of the output file is not the same as the input file or the record format of a record provided by an E15 exit.

System action: Processing stops.

User response: Correct the data set specification and resubmit the job.

CNK208I MIXTURE OF SORTWK DEVICES

Explanation: Sort work data sets were assigned to different device types.

System action: None.

User response: No action is required.

CNK209I xxx PRIMARY AND yyy SECONDARY ddname TRACKS ALLOCATED, zzz USED

Explanation: This message indicates that DB2 Sort needed to request one or more secondary allocations for SORTOUT. xxx is the number of tracks that were initially allocated, yyy is the total number of tracks acquired via secondary allocation, and zzz is the total number of tracks actually required to contain the SORTOUT data set.

System action: None.

User response: No action is required.

**CNK210I E15 RETURN CODE INVALID,
IGNORED**

Explanation: The input exit routine passed a return code of 0 or 4. These return codes are not valid. DB2 Sort ignored the return code and presumed a return code of 8.

System action: None.

User response: Ask the system programmer who is responsible for installing DB2 Sort to contact IBM Software Support.

**CNK219E DYNALLOC FAILED, RETURN
CODE=(nnnn) - uuuuuuuu [-SMS
RETURN CODE=ssss]**

Explanation: The execution of the DYNALLOC macro instruction, which is used to allocate sort work data sets, failed. nnnn indicates the error reason code. Two possible reason codes are: 021C - Undefined unit name. 0214 - Unit not available. If all specified units are unavailable when DYNALLOC is issued, the DYNALLOC request fails. For other reason codes, see z/OS MVS Programming: Authorized Assembler Services Guide, SA22-7608. uuuuuuuu indicates either the unit name or storage class name. ssss indicates the SMS return code. (ssss is displayed only for certain failures that are detected by SMS.)

System action: Processing stops.

User response: If you have specified a DYNALLOC parameter in the DB2SORTI data set incorrectly, correct the parameter and resubmit the job. If the parameter is specified correctly or if you have not specified the DYNALLOC parameter, ask the system programmer who is responsible for installing DB2 Sort to check that the DYNALOC installation option is set correctly.

**CNK219I DYNALLOC FAILED, RETURN
CODE=(nnnn) - uuuuuuuu [-SMS
RETURN CODE=ssss], PROCESSING
CONTINUES**

Explanation: Dynamic allocation was unsuccessful. nnnn indicates the error reason code. uuuuuuuu indicates either the unit name or the storage class name. ssss indicates the SMS return code. (ssss is displayed only for certain failures that are detected by SMS.) Processing continues with previously allocated sort works. For more information about error reason codes, see z/OS MVS Programming: Authorized Assembler Services Guide, SA22-7608

System action: None.

User response: No action is required.

CNK224E ddname NOT DEFINED

Explanation: A required DD statement could not be found.

System action: Processing stops.

User response: Supply the DD statement and resubmit the job.

**CNK227E ddname BLKSIZE GREATER THAN
ASCII LIMIT**

Explanation: The BLKSIZE of an output data set that resides on an ASCII-labeled tape is greater than 2048 bytes; 2048 is the maximum value allowed.

System action: Processing stops.

User response: Correct the BLKSIZE specification and resubmit the job.

**CNK228E ddname DCB BLKSIZE GREATER
THAN TRACK CAPACITY**

Explanation: The BLKSIZE for the indicated output file was greater than the track capacity of the output device.

System action: Processing stops.

User response: Correct the BLKSIZE specification and resubmit the job.

CNK229E ddname DSORG NOT PS/PO

Explanation: The file defined by ddname must be a sequential data set (PS) or a partitioned data set (PO) member.

System action: Processing stops.

User response: Correct the data set specification and resubmit the job.

**CNK232E ddname RECFM=VBS, LRECL
MISSING**

Explanation: A RECFM of VBS was specified for the input ddname without an accompanying LRECL specification.

System action: Processing stops.

User response: Specify the LRECL for the input data set and resubmit the job.

**CNK240E ddname UNSUPPORTED DCB
FUNCTION**

Explanation: The DD statement specified or implied an attribute which is not supported, e.g., hardware keys for a disk output data set or a block prefix length other than 0, 4 or L for an ASCII tape output data set.

System action: Processing stops.

CNK246I • CNK264E

User response: Correct the data set specification and resubmit the job.

CNK246I FILESIZE x

Explanation: x indicates the number of bytes of input data that were sorted or merged by DB2 Sort.

System action: None.

User response: No action is required.

CNK254E ddname VSAM {OPEN,CLOSE} ERROR - xx

Explanation: An error occurred during an attempt to OPEN or CLOSE a VSAM file defined by ddname. xx is the error number. For more information about the error code see DFSMS Macro Instructions for Data Sets, SC26-7408.

System action: Processing stops.

User response: Correct the data set specification and resubmit the job.

CNK255E VSAM LOGICAL ERROR xx ON {INPUT,OUTPUT}

Explanation: An error occurred while processing a VSAM data set. xx is the error number. For more information about the error code see DFSMS Macro Instructions for Data Sets, SC26-7408.

System action: Processing stops.

User response: Correct the VSAM data set error and resubmit the job.

CNK256I ddname VSAM file, RECORDS PADDED ON OUTPUT

Explanation: The fixed-length VSAM LRECL for the indicated output file is greater than the length of the records at the end of DB2 Sort processing. DB2 Sort padded the output records with filler characters on the right.

System action: None.

User response: No action is required.

CNK260I RECOVERY FROM B37 SUCCESSFUL. SORT PROCESSING CONTINUES

Explanation: DB2 Sort recovered from a B37 abend and continued processing.

System action: None.

User response: No action is required.

CNK261E UNEQUAL MAINTENANCE LEVELS: xxxxxxxx,yy,zz

Explanation: The maintenance level of one of the DB2 Sort load modules does not match the maintenance level of the DB2 Sort root module. xxxxxxxx is the name of the load module. yy represents the maintenance level of the xxxxxxxx module. zz represents the maintenance level of the root module.

System action: Processing stops.

User response: Contact the system programmer who is responsible for installing DB2 Sort.

CNK262I REENTRANT SORT NOT RESIDENT - INEFFICIENT SORT

Explanation: The resident DB2 Sort load modules were loaded into the private area instead of being executed from the Link Pack Area or Extended Link Pack Area. This situation might have occurred because the modules were found in a STEPLIB or JOBLIB DD data set. Loading the resident modules into the private area limits the amount of virtual storage that is available to DB2 Sort and can reduce the efficiency of the sort.

System action: None.

User response: Contact the system programmer who is responsible for installing DB2 Sort.

CNK263E MULTIVOLUME SORTWK DATA SETS ARE NOT SUPPORTED

Explanation: DB2 Sort does not support the use of multivolume disk sort work data sets. However, if DB2 Sort requires the use of the space only on the first volume of a multivolume sort work file, this error message will not be issued.

System action: Processing stops.

User response: Remove the multivolume allocation and resubmit the job.

CNK264E UNEQUAL RECORD LENGTHS - VSAM SORTIN - TYPE=F

Explanation: A record in a fixed-length VSAM input data set was encountered whose length was not equal to the length specified in the VSAM cluster definition.

System action: Processing stops.

User response: Correct the record length and resubmit the job.

**CNK265E ddname VSAM CONCATENATED
INPUT NOT ALLOWED**

Explanation: An input file consists of concatenated VSAM data sets. DB2 Sort does not support concatenated VSAM input files.

System action: Processing stops.

User response: Correct the input data set specification and resubmit the job.

**CNK276I SYSDIAG=nnnnnnnn,nnnnn
nnn,nnnnnnnn,nnnnnnnn**

Explanation: This message contains internal diagnostic information intended for use by IBM Software Support.

System action: None.

User response: No action is required.

**CNK300E SCNKLPA LIBRARY NOT
ACCESSIBLE – MODULE CNKSSVC
COULD NOT BE LOADED**

Explanation: The SCNKLPA library must be in the LPA or the LNKLST or accessible through a STEPLIB or a JOBLIB DD statement.

System action: Processing stops.

User response: Ask the system programmer who is responsible for installing DB2 Sort to put the SCNKLPA library either in the LPA or the LNKLST to enable access from the utilities. Alternatively, specify a JOBLIB or STEPLIB DD statement with the name of the SCNKLPA library.

**CNK301E DB2SORTI CONTROL STATEMENT
ERROR – UNKNOWN PARAMETER**

Explanation: A control statement that was read from the DB2SORTI DD data set has a parameter that is not valid.

System action: Processing stops.

User response: Correct the control statement and resubmit the job.

**CNK303E DB2SORTI DYNALLOC PARAMETER
SYNTAX ERROR**

Explanation: The **DYNALLOC** parameter on a control statement that was read from the DB2SORTI data set contains a syntax error.

System action: Processing stops.

User response: Correct the control statement and resubmit the job.

CNK304E DB2SORTI OPEN FAILURE

Explanation: The DB2SORTI data set could not be opened.

System action: Processing stops.

User response: Correct the DB2SORTI data set specification and resubmit the job.

CNK305E DB2SORTO OPEN FAILURE

Explanation: The DB2SORTO data set could not be opened.

System action: Processing stops.

User response: Correct the DB2SORTO data set specification and resubmit the job.

**CNK306E INSUFFICIENT BELOW-THE-LINE
STORAGE**

Explanation: The amount of virtual storage below the 16-MB line is too low for DB2 Sort to continue processing.

System action: Processing stops.

User response: Increase the region size and resubmit the job.

**CNK307E INSUFFICIENT ABOVE-THE-LINE
STORAGE**

Explanation: The amount of virtual storage above the 16-MB line is too low for DB2 Sort to continue processing.

System action: Processing stops.

User response: Increase the region size and resubmit the job.

**CNK310E INTERNAL ERROR: INVALID
PARAMETER LIST LENGTH**

Explanation: The parameter list that was passed to DB2 Sort from the caller is not valid.

System action: Processing stops.

User response: Ask the system programmer who is responsible for installing DB2 Sort to contact IBM Software Support.

**CNK311E INTERNAL ERROR: INVALID
PRODUCT IDENTIFIER**

Explanation: The parameter list that was passed to DB2 Sort from the caller is not valid.

System action: Processing stops.

User response: Ask the system programmer who is

CNK400E • CNK418I

responsible for installing DB2 Sort to contact IBM Software Support.

CNK400E ddname IS AN UNINITIALIZED SEQUENTIAL DISK DATA SET

Explanation: The input data set was allocated but never opened for output. Therefore, there is no valid data or end-of-file mark in the data set. This condition usually occurs when a program abends and the steps to create the data are bypassed.

System action: Processing stops.

User response: Write the appropriate data or end-of-file mark in the data set and resubmit the job.

CNK410I xxx BYTES OF VIRTUAL STORAGE AVAILABLE ABOVE THE 16-MEGABYTE LINE, yyy BYTES RESERVE REQUESTED, zzz BYTES USED

Explanation: This message provides information about the amount of virtual storage that was used. xxx indicates the amount of virtual storage above the 16-MB line that was available when DB2 Sort received control. yyy indicates the amount of virtual storage that the user requested DB2 Sort to reserve above the 16-MB line. zzz indicates the amount of virtual storage that was used by DB2 Sort above the 16-MB line.

System action: None.

User response: No action is required.

CNK411I nnn BYTES OF EMERGENCY SPACE ALLOCATED ABOVE THE 16-MEGABYTE LINE

Explanation: nnn indicates the amount of virtual storage above the 16-MB line that DB2 Sort has reserved for use by other programs (for example, the caller, system SVCs, or the tape management system).

System action: None.

User response: No action is required.

CNK416I {access-method WAS USED FOR ddname} {ddname: EXCP'S=eee [,UNIT=uuuu][,DEV=dddd] [,CHP=ccccccc,n][,VOL=vvvvvv]} {TOTAL OF xxx EXCP'S ISSUED FOR totalid}

Explanation: This message provides summary I/O tuning information for input, output or sort work files that were processed by DB2 Sort. The first form is used when an access method other than EXCP is used for a file. It uses a generic term for the access method, such as BSAM or HIPERBATCH, and the file for which it was used. When EXCP is used, the message takes on

the second form which has the component parts listed below. Some of these components may or may not be included in the message depending on the level of the operating system and the availability of the information within DB2 Sort.

EXCP'S=eee

eee indicates the number of EXCPs that were issued for the file.

UNIT=uuuu

uuuu indicates the unit type on which the data set resides.

DEV=dddd

dddd indicates the device name for the file.

CHP=ccccccc,n

ccccccc indicates the channel paths that are available to the device. n is the number of PAV aliases that are available.

VOL=vvvvvv

vvvvvv identifies the volume serial number of the volume for the file.

For certain types of sorts, DB2 Sort can dynamically allocate data sets other than sort work data sets for use in the sorting process. When non-sort work data sets are used, they are collectively represented in a single CNK416I message using a ddname of "SORTWK&&" for the purpose of reporting EXCPs that were issued against them. In the third form of the message, xxx provides a total of the EXCPs that were issued for SORTWORKS or SORTING, as identified by "totalid".

System action: None.

User response: No action is required.

CNK417I Diagnostic text.

Explanation: Diagnostic message for use by the IBM Support Center.

System action: None.

User response: None.

CNK418I DATASPACE(S) AND/OR ZSPACE USED

Explanation: DB2 Sort has dynamically chosen to use data space or ZSPACE, or a combination of these resources during the execution of the sort. ZSPACE is a replacement for hiperspace. It allows native use of the central storage resources that are available. The ZSPACE technique eliminates the additional overhead that is produced when hiperspace is simulated by the operating system in a z/Architecture® environment. It provides superior CPU performance and reduced system overhead compared to a conventional hiperspace application.

System action: None.

User response: No action is required.

CNK419I Diagnostic text.

Explanation: Diagnostic message for use by the IBM Support Center.

System action: None.

User response: None.

CNK422E SORTOUT STOW FAILURE

Explanation: When writing to SORTOUT, DB2 Sort was unable to store (STOW) the SORTOUT PDS member in the SORTOUT directory. This condition is caused by specifying an insufficient number of directory blocks when creating the SORTOUT data set.

System action: Processing stops.

User response: Increase the number of directory blocks for the output data set and resubmit the job.

CNK423I DYNAMIC ALLOCATION RETRY - WAITING FOR SPACE

Explanation: The DYNALLOC facility is being used to acquire sort work space, but the system currently does not have enough disk space to satisfy the request. DB2 Sort will wait the prescribed number of minutes as specified by the **RETRY** parameter and then try the request again.

System action: None.

User response: No action is required.

CNK424I DYNAMIC ALLOCATION RETRY SUCCESSFUL

Explanation: The dynamic allocation of sort work space after a DYNALLOC RETRY attempt was successful. Sort processing continues.

System action: None.

User response: No action is required.

CNK428I CALLER-PROVIDED IDENTIFIER IS "xxxx"

Explanation: DB2 Sort was called using a 31-bit parameter list where the **call identifier** parameter was specified. xxxx is the identifier that was specified by the caller.

System action: None.

User response: No action is required.

CNK432I {SORT,MERGE} FORMAT OPERAND IGNORED

Explanation: On either a SORT or MERGE control statement, the format of the keys was specified in both the **FIELDS** and **FORMAT** parameters. DB2 Sort ignores the **FORMAT** parameter and uses the individual format specifications within the **FIELD** parameter.

System action: None.

User response: No action is required.

CNK447I PHASE 3 VIRTUAL STORAGE REDUCED TO nnn BYTES FOR OPTIMAL PERFORMANCE

Explanation: Phase 3 optimization has determined that a reduction in the amount of virtual storage is appropriate for an efficient execution. nnn is the amount of virtual storage used during phase 3. See message "CNK164I" on page 61 to determine the amount of virtual storage that was used during earlier phases of the sort execution.

System action: None.

User response: No action is required.

CNK460I ddname DATA TRUNCATED DUE TO DCB BLKSIZE OVERRIDE

Explanation: An extended sequential data set used as input has had its DCB BLKSIZE overridden to a smaller value via a JCL specification. A physical block exceeding this overridden BLKSIZE specification was truncated to the smaller size during input processing.

System action: None.

User response: Confirm that this truncation is desired. If not, remove the BLKSIZE specification from the JCL before resubmitting the job.

CNK462I OUTPUT LRECL DIFFERS FROM SORTOUT LRECL

Explanation: The LRECL defined in the JCL differs from the SORTIN LRECL or the internally processed record length. Fixed-length records will be padded to the SORTOUT LRECL when the SORTOUT LRECL is greater than the SORTIN or internally processed record length. Records will be truncated to the SORTOUT LRECL when the SORTOUT LRECL is less than the SORTIN or internally processed record length.

System action: None.

User response: Verify that the padding or truncation that will be performed is desired for this application. Refer to the CNK108I and CNK110I messages that provide the input and output record lengths. If the padding or truncation is not desired, modify the LRECL specification before resubmitting the job.

CNK463E ddname IS A LINEAR VSAM DATA SET

Explanation: DB2 Sort does not support an input or output file that is a linear VSAM data set.

System action: Processing stops.

User response: Correct the data set specification and resubmit the job.

Explanation: An unexpected internal error occurred.

System action: Processing stops.

User response: Ask the system programmer who is responsible for installing DB2 Sort to contact IBM Software Support.

CNK493I ZIIP PROCESSOR USED

Explanation: DB2 Sort used the zIIP processor for improved performance.

System action: None.

User response: No action is required.

CNK555I OPTMODE=x

Explanation: This message displays the optimization mode in effect. x is BALANCE, CPU or ELAP. The optimization mode is set at installation time, but the OPTMODE runtime parameter can be used to override the installation default.

System action: None.

User response: No action is required.

CNK494I {INPUT, OUTPUT} PHASE USED MIDAW {;MIXED MODE}

Explanation: This message indicates that DB2 Sort MIDAW technology optimized the performance of the input and/or output phase. If MIXED MODE is displayed in the message, the device allocations contained a mixture of MIDAW-capable and MIDAW-incapable devices. Performance might not be optimal in this situation.

System action: None.

User response: No action is required.

CNK998E jobname ,stepname, procstepname – UNSUCCESSFUL SORT nnnt

Explanation: A critical error occurred. nnn is the hexadecimal abend or message number. t is the error type. S indicates that the error is a system abend, U indicates that the error is a user abend, and W indicates that an internal error occurred.

System action: Processing stops.

User response: Contact the system programmer who is responsible for installing DB2 Sort.

Programmer response: Contact IBM Software Support.

CNK526E DB2 SORT FOR Z/OS INTERNAL ABEND

Gathering diagnostic information

Before you report a problem with DB2 Sort to IBM Software Support, you need to gather the appropriate diagnostic information.

Procedure

Provide the following information for all DB2 Sort problems:

- A clear description of the problem and the steps that are required to re-create the problem
- All messages that were issued as a result of the problem
- Product release number and the number of the last program temporary fix (PTF) that was installed
- The version of DB2 that you are using and the type and version of the operating system that you are using

Tip: If you can rerun the problem step, add a DB2SORTI DD statement followed by the DEBUG parameter, for example:

```
//DB2SORTI DD *
DEBUG
```

Then rerun the failing step. The **DEBUG** parameter will provide additional diagnostic information.

Chapter 6. Tools Customizer reference

Before you use Tools Customizer, you should understand the Tools Customizer terminology and the data sets that Tools Customizer uses during customization.

Tools Customizer terminology and data sets

Before you use Tools Customizer, you should understand the Tools Customizer terminology and the data sets that Tools Customizer uses during customization.

Tools Customizer terminology

Tools Customizer uses several unique terms that you should be familiar with before you begin to use Tools Customizer.

Products and components

How an IBM Tool is packaged determines whether it is referred to as a product or as a component in the Tools Customizer documentation and interface. An IBM Tool that is ordered as a stand-alone entity (that is, not as part of a solution pack) is referred to as a product. An IBM Tool that is part of a solution pack is referred to as a component. Some IBM Tools are available in both formats; therefore, the same IBM Tool can be referred to as a product or as a component depending on how it is packaged.

DB2 entry

You can customize DB2 Sort on one or more DB2 entries. A DB2 entry can be any of the following items:

DB2 subsystem

A distinct instance of a relational database management system (RDBMS) that is not part of a data sharing group. An example of a DB2 subsystem name is DB01.

DB2 group attach name

The name that is used by the TSO/batch attachment, the call attachment facility (CAF), DL/I batch, utilities, and the Resource Recovery Services attachment facility (RRSAF) as a generic attachment name. An example of a group attach name is DSG1.

DB2 data sharing member

A DB2 subsystem that is assigned by the cross-system coupling facility (XCF) to a data sharing group. An example of a DB2 data sharing member name is DB02.

Tools Customizer maintains the following lists of DB2 entries:

Associated list

The list of DB2 entries that are associated with DB2 Sort. If the product to be customized requires DB2 entries, you can customize DB2 Sort only on DB2 entries that are in the associated list. When you customize DB2 Sort, this list is displayed in the DB2 Entries, Associations, and Parameter Status section of the Customizer Workplace panel.

You can add and copy DB2 entries to the associated list. When you add or copy DB2 entries to the associated list, the entries are associated with DB2 Sort.

Master list

The list of all DB2 entries that are defined but are not associated with DB2 Sort. Tools Customizer obtains information about these DB2 entries either from entries that were created manually or from the customizations of other products that were discovered. If you remove a DB2 entry from the associated list, the DB2 entry is added to the master list. When you create a new DB2 entry, it is added to the master list, and when you associate the new entry with DB2 Sort, it is removed from the master list and added to the associated list. The master list is displayed on the Associate a DB2 Entry for Product panel.

If the associated list does not have the DB2 entries on which you want to customize DB2 Sort, you can associate existing entries from the master list to the associated list.

You can create new DB2 entries and copy existing entries to the master list.

High-level qualifier

The high-level qualifier is considered to be all of the qualifiers except the lowest level qualifier. A high-level qualifier includes a mid-level qualifier.

Product parameters

Parameters that are specific to DB2 Sort. These parameters are defined by DB2 Sort and are stored in a data member that is defined by DB2 Sort.

DB2 parameters

Parameters for a DB2 entry. These parameters are defined by Tools Customizer and are stored in a DB2 parameter data member.

Status type

Product, LPAR, and DB2 entry status type

After you specify the product that you want to customize, the product, the LPAR, and the DB2 entries have a status. The status is partly based on whether required parameters are defined. For some products, LPAR parameters or DB2 parameters might not be required. In these cases, the status is Not Required.

To customize DB2 Sort, all of the required parameters must be defined.

If required parameters for the the product parameters or DB2 parameters are not defined, the status of the parameters is Incomplete. Define values for parameters by manually editing them or by generating the customization jobs and specifying values for all of the required parameters that are displayed on the panels.

When values for all of the required parameters are defined, the status is Ready to Customize. Customization jobs can be generated only when all of the required parameters are defined and the status is Ready to Customize or Customized for the product parameters and DB2 parameters for the DB2 entries on which DB2 Sort will be customized.

The following table shows the meaning of the status types. Each status is defined differently for each type of parameter.

Table 13. Status types for the product, the LPAR, and the DB2 entries

Status	Product	LPAR	DB2 entries
Incomplete	The required product parameters are not defined, or the required product parameters are defined but LPAR parameters, DB2 parameters, or both are not defined.	The required parameters are not defined.	The required parameters are not defined.
Discovered	The product parameter definitions were discovered by using the product Discover EXEC.	N/A	N/A
Ready to Customize	The required product, LPAR, and DB2 parameters are defined, the status is Ready to Customize or Customized for the LPAR and at least one associated DB2 entry. You can generate the customization jobs.	The required LPAR parameters are defined or LPAR parameters are not required.	The required DB2 parameters are defined or DB2 parameters are not required.
Customized	The jobs are customized on the local LPAR.	The jobs are customized for the product or for all of the associated DB2 entries on the local LPAR.	The jobs are customized for the DB2 entry.
Errors in Customization	N/A	N/A	Errors occurred while the customization jobs were being generated.
Not Required	N/A	LPAR parameters are not required.	DB2 parameters are not required.

Related tasks:

“Creating and associating DB2 entries” on page 40

DB2 entries are optional for DB2 Sort. You can create new DB2 entries and associate them with DB2 Sort.

“Copying DB2 entries” on page 47

You can copy associated and not associated DB2 entries to other DB2 entries or to new DB2 entries.

“Removing DB2 entries” on page 49

You can remove DB2 entries from the associated list.

Data sets that Tools Customizer uses during customization

Tools Customizer uses several unique data sets during the customization process. Familiarize yourself with these data sets before you begin to use Tools Customizer.

Several different data sets are required to customize DB2 Sort with Tools Customizer. These data sets are supplied by DB2 Sort, supplied by Tools Customizer, or allocated by Tools Customizer.

DB2 Sort provides the following data sets:

Metadata library

Contains the metadata for the product to be customized. Tools Customizer uses the metadata to determine which tasks, steps, and parameters to display on the Product Parameters panel, the LPAR Parameters panel, and the DB2 Parameters panel. This data set also contains the templates that Tools Customizer uses to generate the customization jobs.

The metadata library naming convention is *high_level_qualifier*.SCNKDENU, where *high_level_qualifier* is all of the segments of the data set name except the lowest-level qualifier.

You specify the metadata library on the Specify the Metadata Library panel. READ access to this data set is required.

Discover EXEC library

Contains the DB2 Sort Discover EXEC. When you customize DB2 Sort, you can use the Discover EXEC to automatically retrieve and store product information, such as parameter values from an already customized product. Tools Customizer saves the discovered information in the data store.

The default name of the data set is the high-level qualifier for the metadata library plus a lowest-level qualifier. For DB2 Sort, the lowest-level qualifier is SCNKDENU. You can change the default value on the Discover Customized Product Information panel. EXECUTE access to this data set is required.

Tools Customizer provides the following data sets:

Tools Customizer metadata library

Contains the metadata for the DB2 and LPAR parameters that are required to customize DB2 Sort. Tools Customizer uses the metadata to determine which parameters to display on the DB2 Parameters panel and the LPAR Parameters panel. In addition, Tools Customizer uses information in the metadata library to determine whether additional DB2 and LPAR parameters need to be displayed on these panels. As you customize different products, different DB2 and LPAR parameters might need to be defined.

The default name of the data set is DB2TOOL.CCQ110.SCCQDENU. You can change the default value on the Tools Customizer Settings panel. READ access to this data set is required.

Tools Customizer table library

Stores information about jobs that are customized. Job information that is stored includes a description of the job, its member name and template name, the SSID, group attach name, and when the job was generated.

The default name of the data set is DB2TOOL.CCQ110.SCCQTENU. WRITE access to this data set is required.

Tools Customizer requires that the following data sets exist during the customization process. If the data sets do not exist, Tools Customizer automatically allocates them.

Discover output data set

Contains the output that is generated when you run the DB2 Sort Discover EXEC. The DB2 Sort Discover EXEC retrieves the metadata and values for the parameters from a previous customization of DB2 Sort.

The default name of the data set is DB2TOOL.CCQ110.DISCOVER. You can change the default value on the Tools Customizer Settings panel or the Discover Customized Product Information panel. WRITE access to this data set is required.

Data store data set

Contains product, LPAR, and DB2 parameter values, and DB2 entry associations. Tools Customizer uses this data set to permanently store all information that is acquired about the product, DB2 subsystems or data sharing groups, and LPAR when you customize products on the local LPAR.

The default name of the data set is DB2TOOL.CCQ110.DATASTOR. You can change the default value on the Tools Customizer Settings panel. WRITE access to this data set is required.

Customization library

Contains the customization jobs that Tools Customizer generates for DB2 Sort.

Tools Customizer checks whether a customization library name was specified for more than one instance of the same version of the same product. If the same customization library name is specified for more than one product of the same version, the CCQD123E message is issued to prevent you from overwriting previously generated customization jobs. Ensure that you specify unique qualifier for the customization library for each instance of the product.

To customize DB2 Sort, submit the members of the data set in the order in which they are displayed on the Finish Product Customization panel.

The data set naming convention is *hlq.\$LPAR_name\$.xyzvrm*, where:

- *hlq* is the value of the **Customization library qualifier** field on the Tools Customizer Settings panel (CCQPSET)
- *LPAR_name* is the four-character LPAR name
- *xyzvrm* is the three-letter product identifier with the version, release, and modification level

For example, the data set name might be DB2TOOL.PRODUCT.CUST.\$MVS1\$.XYZ410.

WRITE access to this data set is required.

Tools Customizer allocates the data sets for the discover output, the data store, and the customization library with the attributes that are shown in the following table:

Table 14. Data set attributes for allocating the Discover output, data store, and customization library data sets

Data set	Organization	Record format	Record length	Block size	Data set name type
Discover output data set	PO	Variable block	16383	32760	LIBRARY

Table 14. Data set attributes for allocating the Discover output, data store, and customization library data sets (continued)

Data set	Organization	Record format	Record length	Block size	Data set name type
Data store data set	PO	Variable block	16383	32760	LIBRARY
Product customization library	PO	Fixed block	80	32720	LIBRARY

Restrictions:

- Multiple users cannot simultaneously share the discover output data set, data store data set, Tools Customizer metadata library, and metadata library.
- You cannot share the data store data set across multiple LPARs with shared DASD or copy the data store data set to another LPAR. Tools Customizer creates many cross-references between product and DB2 associations. Therefore, if you share or copy the data store data set, member names that are empty or that do not exist might be generated.

Notices

This information was developed for products and services offered in the U.S.A.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

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

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

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

IBM Corporation
J64A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

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

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

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

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

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

COPYRIGHT LICENSE:

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

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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

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

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

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.html>.

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

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and the section titled "Cookies, Web Beacons, and Other Technologies" in IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details>. Also, see the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Index

A

- accessibility
 - overview 8
- accessing DB2 Sort libraries 52
- activating DB2 Sort 27

C

- changing display options 32
- components 6
- cookie policy 79, 81
- Copy DB2 Entries panel 47
- customization
 - associated list
 - adding DB2 entries 40
 - overview 18, 73
 - associating DB2 entries 40
 - browsing parameters 47
 - changing display options 32
 - changing parameters 33
 - component 18, 73
 - copying DB2 entries 47
 - Create a DB2 Entry panel 40
 - creating DB2 entries 40
 - customization jobs
 - deleting 50
 - displaying 50
 - generating 45
 - maintaining 50
 - regenerating 45
 - renaming 50
 - sort sequence 45
 - submitting 45, 50
 - customization library
 - deleting jobs 50
 - maintaining 50
 - overview 21, 76
 - recustomizing 50
 - renaming jobs 50
 - customization library qualifier
 - specifying 29
 - Customized status 18, 73
 - Customizer Workplace panel 45
 - customizing a new version of a product 33
 - customizing a product for the first time 33
 - customizing settings 29
- data sets
 - customization library 21, 76
 - data store 21, 76
 - Discover EXEC library 21, 76
 - metadata library 21, 76
- data store
 - overview 21, 76
- data store data set
 - specifying 29
- DB2 data sharing members
 - adding 40
 - associating 40

- customization (*continued*)
 - DB2 data sharing members (*continued*)
 - copying 47
 - creating 40
 - DB2 entries 18, 73
 - adding 40
 - associating 40
 - copying 47
 - creating 40
 - defining 45
 - deleting 49
 - generating jobs for 45
 - removing 49
 - selecting 45
 - specifying 45
 - unassociating 49
 - DB2 group attach field
 - specifying 29
 - DB2 group attach names
 - adding 40
 - associating 40
 - copying 47
 - creating 40
 - DB2 parameters
 - defining 43
 - editing 43
 - DB2 Parameters panel 43
 - DB2 subsystems
 - adding 40
 - associating 40
 - copying 47
 - creating 40
 - defining DB2 parameters 43
 - defining parameters 41, 45
 - defining product parameters 42
 - deleting DB2 entries 49
 - deleting jobs 36
 - Discover Customized Product Information panel 38
 - Discover EXEC
 - customizing a new version of a product 33, 34
 - overview 21, 76
 - retrieving product information automatically 38
 - Discovered status 18, 73
 - discovering previous versions 34
 - discovering product information 38
 - display options 32
 - displaying jobs 50
 - displaying panel text 32
 - editing parameters 33
 - editing product parameters 42
 - Errors in Customization status 18, 73
 - Finish Product Customization panel 45
 - first-time 33
 - first-time customization 33
 - generating jobs 45
 - hiding panel text 32
 - high-level qualifier 18, 73

- customization (*continued*)
 - Incomplete status 18, 73
 - job sort order 45
 - jobs
 - deleting 50
 - displaying 50
 - maintaining 50
 - renaming 50
 - sort order 45
 - submitting 45, 50
 - LPARs 51
 - maintaining jobs 50
 - master list
 - adding DB2 entries 40
 - Associate DB2 Entry for Product panel 40
 - overview 18, 73
 - maximizing information panels 32
 - metadata libraries
 - specifying 37
 - metadata library
 - overview 21, 76
 - specifying 29
 - modifying parameters 33
 - modifying settings 29
 - multiple instances 29
 - multiple-LPAR environment 51
 - Not Required status 18, 73
 - options 32
 - panel display options 32
 - panels
 - Associate DB2 Entry for Product 40
 - Create a DB2 Entry 40
 - Customizer Workplace 45
 - DB2 Parameters 43
 - Discover Customized Product Information 38
 - Finish Product Customization 45
 - Product Parameters 42
 - Specify the Metadata Library 37
 - parameters
 - browsing 47
 - defining 41, 45
 - viewing 47
 - preparing to use Tools Customizer 29
 - product 18, 73
 - product parameters
 - changing 36
 - defining 42
 - editing 36, 42
 - modifying 36
 - Product Parameters panel 42
 - Ready to Customize status 18, 73
 - recustomization 33, 36
 - recustomizing 36
 - recustomizing a product 33
 - removing DB2 entries 49
 - roadmaps 33
 - customizing for the first time 33

- customization (*continued*)
 - roadmaps (*continued*)
 - first-time customization 33
 - reustomizing 36
 - Specify the Metadata Library panel 37
 - specifying data sets 29
 - specifying metadata libraries 37
 - starting Tools Customizer 28
 - status types
 - Customized 18, 73
 - Discovered 18, 73
 - Errors in Customization 18, 73
 - Incomplete 18, 73
 - Not Required 18, 73
 - Ready to Customize 18, 73
 - submitting jobs 45
 - terminology 18, 73
 - user job card settings
 - specifying 29
 - viewing parameters 47
- customization library
 - overview 21, 76
- customization library qualifier
 - specifying 29
- customizing settings 29

D

- data store
 - overview 21, 76
- data store data set
 - specifying 29
- DB2 group attach field
 - specifying 29
- DB2 Sort - accessing libraries 52
- DB2 Sort - activating 27
- DB2 Sort - scenarios 6
- DB2 Sort libraries
 - SCNKLINK 52
 - SCNKLPA 52
- DB2SORT subsystem parameter 23, 27
- DB2SORTI 23, 70
- DEBUG parameter 70
- default installation options 16
- diagnostic information
 - gathering 70
- Discover EXEC
 - overview 21, 76
- display options 32
- displaying panel text 32
- documentation
 - accessing 7
 - sending feedback 7
- documentation changes 8
- DYNALLOC runtime parameter 23
- dynamic sort work allocation
 - installation options 16
 - runtime parameter 23

F

- first-time customization 33
- functions
 - overview 2

H

- hiding panel text 32

I

- Installation Verification Program - IVP 8, 53
- Installation Verification Program (IVP)
 - overview 53
 - report generation 54
 - using 54
- inter-regional parallelism 2, 5, 8
- intra-regional parallelism 2, 5
- IVP - Installation Verification Program 8, 53

L

- legal notices
 - cookie policy 79, 81
 - notices 79
 - programming interface information 79
 - trademarks 79, 81

M

- maximizing information on panels 32
- memory usage optimization 2
- message format 57
- messages
 - overview 57
- metadata library
 - overview 21, 76
 - specifying 29
- modifying settings 29

N

- notices 79

O

- optimization mode option 16
- options 32
- OPTMODE
 - installation option 8, 16
 - runtime parameter 8, 23
- overview 1

P

- PAGEMON
 - installation option 16
- panel display options 32
- panels
 - Copy DB2 Entries 47
- preparing to use Tools Customizer 29
- problems
 - diagnostic information about 70
- process flow overview 5
- programming interface information 79

R

- reader comment form 7
- REGION parameter 2, 23
- roadmaps
 - customizing for the first time 33
 - first-time customization 33
- runtime considerations 23
- runtime parameters
 - DYNALLOC 23
 - OPTMODE 8, 23
 - TUNEMSG 23

S

- scenarios 6
- SCNKLINK - DB2 Sort library 52
- SCNKLPA - DB2 Sort library 52
- screen readers and magnifiers 8
- sensitivity to paging activity option 16
- service information 7
- Sort Capacity Exceeded error 2
- sort processing 2
- sort work space usage 2
- SORTNUM parameter 2, 27
- specifying data sets 29
- summary of changes 8
- support information 7

T

- technotes 7
- Tools Customizer 6
 - associated list
 - adding DB2 entries 40
 - overview 18, 73
 - associating DB2 entries 40
 - browsing parameters 47
 - changing display options 32
 - component 18, 73
 - Copy DB2 Entries panel 47
 - copying DB2 entries 47
 - Create a DB2 Entry panel 40
 - creating DB2 entries 40
 - customization jobs
 - deleting 50
 - displaying 50
 - generating 45
 - maintaining 50
 - renaming 50
 - sort sequence 45
 - submitting 45, 50
 - customization library
 - deleting jobs 50
 - maintaining 50
 - recustomizing 50
 - renaming jobs 50
 - customization library qualifier
 - specifying 29
 - Customized status 18, 73
 - Customizer Workplace panel 45
 - customizing a new version of a product 33, 34
 - customizing a product for the first time 33
 - data sets
 - customization library 21, 76

Tools Customizer (*continued*)

- data sets (*continued*)
 - data store 21, 76
 - Discover EXEC library 21, 76
 - metadata library 21, 76
- data store data set
 - specifying 29
- DB2 data sharing members
 - adding 40
 - associating 40
 - copying 47
 - creating 40
- DB2 entries 18, 73
 - adding 40
 - associating 40
 - copying 47
 - creating 40
 - defining 45
 - deleting 49
 - generating jobs for 45
 - removing 49
 - selecting 45
 - specifying 45
 - unassociating 49
- DB2 group attach field
 - specifying 29
- DB2 group attach names
 - adding 40
 - associating 40
 - copying 47
 - creating 40
- DB2 parameters
 - defining 43
 - editing 43
- DB2 Parameters panel 43
- DB2 subsystems
 - adding 40
 - associating 40
 - copying 47
 - creating 40
- defining DB2 parameters 43
- defining parameters 41, 45
- defining product parameters 42
- deleting DB2 entries 49
- deleting jobs 36
- Discover Customized Product
 - Information panel 38
- Discover EXEC
 - customizing a new version of a product 33, 34
 - retrieving product information automatically 38
- Discovered status 18, 73
- discovering product information 38
- displaying jobs 50
- displaying panel text 32
- editing product parameters 42
- Errors in Customization status 18, 73
- features 6
- Finish Product Customization
 - panel 45
- first-time customization 33
- generating jobs 45
- hiding panel text 32
- high-level qualifier 18, 73
- Incomplete status 18, 73
- job sort order 45

Tools Customizer (*continued*)

- jobs
 - deleting 50
 - displaying 50
 - maintaining 50
 - renaming 50
 - submitting 50
- maintaining jobs 50
- master list
 - adding DB2 entries 40
 - Associate DB2 Entry for Product panel 40
 - overview 18, 73
- maximizing information on panels 32
- metadata libraries 37
 - specifying 37
- metadata library
 - specifying 29
- multiple instances 29
- multiple-LPAR environment 51
- new component 8
- Not Required status 18, 73
- overview 6
- panels
 - Associate DB2 Entry for Product 40
 - Copy DB2 Entries 47
 - Create a DB2 Entry 40
 - Customizer Workplace 45
 - DB2 Parameters 43
 - Discover Customized Product Information 38
 - Finish Product Customization 45
 - Product Parameters 42
 - Specify the Metadata Library 37
- parameters
 - browsing 47
 - viewing 47
- preparing to use 29
- product 18, 73
- product parameters
 - changing 36
 - editing 36
 - modifying 36
- Product Parameters panel 42
- Ready to Customize status 18, 73
- recustomization 33
- recustomizing a product 33, 36
- removing DB2 entries 49
- roadmaps
 - customizing a new version of a product 34
 - recustomizing a product 36
 - using the Discover EXEC 34
- Specify the Metadata Library
 - panel 37
 - specifying metadata libraries 37
- starting 28
- status types
 - Customized 18, 73
 - Discovered 18, 73
 - Errors in Customization 18, 73
 - Incomplete 18, 73
 - Not Required 18, 73
 - Ready to Customize 18, 73
- submitting jobs 45

Tools Customizer (*continued*)

- terminology 18, 73
- user job card settings
 - specifying 29
 - using the Discover EXEC 34
 - viewing parameters 47
- trademarks 79, 81
- troubleshooting 57
- TUNEMSG
 - installation option 16
 - runtime parameter 23
 - tuning messages option 16

U

- user job card settings
 - specifying 29
- UTSORTAL parameter 2, 27

Z

- zIIP technology 2
- zIIP usage 8



Product Number: 5655-AA9

Printed in USA

SC19-4387-00

