

Advanced User's Guide

!DB®/WORKBENCH for DB2

Version 500

Document Number TW54-5848-3

June 2000

**Candle Corporation
201 North Douglas Street
El Segundo, California 90245**

Registered trademarks and service marks of Candle Corporation: AF/OPERATOR, AF/PERFORMER, AF/REMOTE, Availability Command Center, Candle Command Center, Candle Electronic Customer Support, Candle Logo, Candle Management Server, Candle Management Workstation, Candle Technologies, CL/CONFERENCE, CL/SUPERSESSION, CT, CT/Data Server, CT/DS, DB Logo, DB/QUICKCHANGE, DELTAMON, ETEWatch, IntelliWatch, MQSecure, MQView, OMEGACENTER, OMEGAMON, OMEGAMON/e, OMEGAMON II, OMEGAMON Monitoring Agent, OMEGAVIEW, OMEGAVIEW II, Solutions for Networked Businesses, and Transplex.

Trademarks and service marks of Candle Corporation: Alert Adapter, Alert Adapter Plus, Alert Emitter, AMS, Amsys, AUTOMATED FACILITIES, Availability Management Systems, Candle Business Partner Logo, Candle Direct Logo, CandleLight, CandleNet, CandleNet 2000, CandleNet Portal, CCC, CECS, CICAT, CL/ENGINE, CL/GATEWAY, CL/TECHNOLOGY, CMS, CMW, Command & Control, CommandWatch, Connect-Two, CSA ANALYZER, CT/ALS, CT/Application Logic Services, CT/DCS, CT/Distributed Computing Services, CT/Engine, CT/Implementation Services, CT/IX, CT/Workbench, CT/Workstation Server, CT/WS, DB/DASD, DB/EXPLAIN, DB/MIGRATOR, DB/QUICKCOMPARE, DB/SMU, DB/Tools, DB/WORKBENCH, Design Network, DEXAN, eBA*ServiceMonitor, End-to-End, Enterprise Candle Command Center, Enterprise Candle Management Workstation, EPILOG, ERPNet, ESRA, HostBridge, IntelliWatch Pinnacle, Lava Console, Messaging Mastered, MQADMIN, MQEdit, MQEXPERT, MQMON, NBX, OMA, OMC Gateway, OMC Status Manager, OMEGACENTER Bridge, OMEGACENTER Gateway, OMEGACENTER Status Manager, OMEGAMON Management Center, OSM, PC COMPANION, Performance Pac, PowerQ, PQConfiguration, PQEdit, PQScope, Response Time Network, Roma, Roma Broker, Roma BSP, Roma Connector, Roma Developer, Roma FS/A, Roma FS/Access, Roma Network, Roma Systems Manager, Roma Workflow Access, Roma WF/Access, RTA, RTN, SentinelManager, Solutions for Networked Applications, Status Monitor, Tracer, Unified Directory Services and Volcano.

Trademarks and registered trademarks of other companies: AIX, DB2, and MQSeries are registered trademarks of International Business Machines Corporation. SAP is a registered trademark and R/3 is a trademark of SAP AG. UNIX is a registered trademark in the U.S. and other countries, licensed exclusively through X/Open Company Ltd. HU-UX is a trademark of Hewlett-Packard Company. SunOS is a trademark of Sun Microsystems, Inc.

All other company and product names used herein are trademarks or registered trademarks of their respective companies.

ProtoView Development Corp. - May contain DataTable Version 3.0 Copyright 1989—1996 by ProtoView Development Corp. and distributed under license by Candle Corporation.

Copyright © 1992, 2000, Candle Corporation, a California corporation. All rights reserved. International rights secured.

Threaded Environment for AS/400, Patent No. 5,504,898; Data Server with Data Probes Employing Predicate Tests in Rule Statements (Event Driven Sampling), Patent No. 5,615,359; MVS/ESA Message Transport System Using the XCF Coupling Facility, Patent No. 5,754,856; Intelligent Remote Agent for Computer Performance Monitoring, Patent No. 5,781,703; Data Server with Event Driven Sampling, Patent No. 5,809,238; Threaded Environment for Computer Systems Without Native Threading Support, Patent No. 5,835,763; Object Procedure Messaging Facility, Patent No. 5,848,234; Communications on a Network, Patent Pending; End-to-End Response Time Measurement for Computer Programs, Patent No. 5,999,705; Improved Message Queuing Based Network Computing Architecture, Patent Pending; User Interface for System Management Applications, Patent Pending.

NOTICE: This documentation is provided with RESTRICTED RIGHTS. Use, duplication, or disclosure by the Government is subject to restrictions set forth in the applicable license agreement and/or the applicable government rights clause.

This documentation contains confidential, proprietary information of Candle Corporation that is licensed for your internal use only. Any unauthorized use, duplication, or disclosure is unlawful.

Contents

Read This First	13
Preface	13
Conventions	15
Adobe Portable Document Format	16
Documentation Set	17
What's New	19
Version 500 Changes	19
Using SQL and DB2	21
Chapter 1. Using the Dynamic SQL Facility in !DB/WORKBENCH	23
Introduction	23
Chapter contents	23
Resources for Learning about Features Available from !DB/WORKBENCH	24
Overview of the Chapter	25
Accessing the SQL Input Panel	27
Overview of the SQL Input Panel	28
Using Uppercase and Lowercase	30
Guidelines for SQL Statements Entered Directly	31
Executing Statements Entered Directly	32
Executing Retrieved Statements	33
Recovery from Errors	34
Keyboard Mapping and SQL Statements Stored in a PDS Member	35
Guidelines for SQL Statements Stored in a PDS Member ..	36
Executing Statements Stored in a PDS Member	38
Interpreting the Report Displayed by the System	40
Testing SQL Statements within a Program	42
Chapter 2. Using DB2 Commands in !DB/WORKBENCH	45
Introduction	45
Chapter contents	45
Overview of the Chapter	46
Resources for Using DB2 Commands	49
Accessing the DB2 Commands Menu	50
Accessing Options on the DB2 Commands Menu	52
Overview of DB2 Commands and Your Tasks	53
Overview of Operation of the Options and Output	59
Introduction to DB2 Commands for Databases	61
Introduction to the Command to Stop DB2	66
Introduction to DB2 Commands for DDF	67

Introduction to DB2 Commands for RLF	69
Introduction to DB2 Commands for Threads	72
Introduction to DB2 Commands for Traces	77
Introduction to DB2 Commands for Utilities	82
Introduction to DB2 Commands for the Archive Log	83
Introduction to DB2 Commands for Buffer Pools	85
Introduction to DB2 Commands to Recover BSDS	91
Introduction to DB2 Commands for Data Sharing	92
Managing the System	95
Chapter 3. Using Authorizations to Manage the System	97
Introduction	97
Chapter contents	97
Resources for Learning about Topics in This Section	98
Overview of the Chapter	99
Resources for Learning about Authorizations	101
Introduction to Methods to Display Authorizations	102
Accessing the Authorizations Menu	104
Displaying Authorizations from the Menu	105
Displaying Authorizations from an Object List Panel	106
Controlling the Display of Authorizations	107
Tasks Commonly Performed with Authorizations	108
Chapter 4. Using the DASD Estimators	111
Introduction	111
Chapter contents	111
Overview of the Chapter	112
Resources for Using the Estimators	113
Accessing the Estimators	114
Functions Available with the Estimators	116
Prerequisites for Estimating DASD Requirements	118
Values Reported by the Estimator	120
Chapter 5. Using the DSNZPARMs Facility	123
Introduction	123
Chapter contents	123
Overview of the Chapter	124
Functions Available with the DSNZPARMs Facility	126
Accessing the DSNZPARMs Menu	128
Selecting Options from the DSNZPARMs Menu	130
Displaying and Editing the DSNZPARMs Parameters	131
Generating Macros from the DSNZPARMs Menu	134
Migrating DB2 Parameters	136
Printing DB2 Parameters	137
Chapter 6. !DB/WORKBENCH Housekeeping Options	139

Introduction	139
Chapter contents	139
Overview of the Chapter	140
Introduction to !DB/WORKBENCH Housekeeping Options	143
Overview of !DB/WORKBENCH Housekeeping	144
Accessing !DB/WORKBENCH Housekeeping Options	146
Option 1: Profile Data Administration	148
Option 1.0: Global Dataset Information	150
Option 1.1: Global Configuration Information	153
Option 1.2: Global Print Options	156
Option 1.3: DB2 Subsystem-Configuration Information	159
Option 1.4: Catalog Prefixes for the IDF	163
Option 1.5: !DB/WORKBENCH Configuration Information	165
Option 1.6: !DB/WORKBENCH Processing Information	169
Option 2: Update Authorizations Information	172
Option 3: Display Extract History	173
Option 4: Run a New Extract	178
Option 5: Display System PDS Members	180
Option 6: Resurrect Old Catalog Copies	182
Option 7: Display the Activity Log	186
Option 8: Display the !DB/Tools Allocated Profile Datasets	189
Option 9: Set the !DB/Tools Profile Datasets	191
Chapter 7. Managing Statistics in the STATDBAS Table Space	195
Introduction	195
Chapter contents	195
Overview of the Chapter	196
General Background about the STATDBAS Table Space	198
Background about Growth in Size of Tables	201
Controlling the Size of Tables	203
Contents of the Tables in STATDBAS Table Space	204
Creating Objects	207
Chapter 8. Introduction to Creating Objects	209
Introduction	209
Chapter contents	209
Resources for Learning about Topics in This Section	210
Overview of Chapter	212
Overview of Methods to Create Objects	214
How to Access and Select the Option You Want	216
How to Determine Order of Creation of Objects	218
Methods Available If You Make a Mistake	219
Introduction to Methods for Providing Values	220
Methods Available to Estimate Requirements for Space	222
Overview of Operation of Defaults	223
Background about Memory and Options	226
Memory and Effects on Selecting from Lists	227

Memory and Effects on Ending Sessions	228
Memory and Effects on Beginning a Session	230
Processing the Statements	232
Chapter 9. Creating a Database	235
Introduction	235
Chapter contents	235
Overview of Chapter	236
Access to the Create Database Panel	237
Prerequisites for Creating a Database	238
Beginning to Create the Database	239
Using Values from an Existing Database	240
Specifying an Existing Storage Group for the New Database	242
Creating a New Storage Group for the New Database	244
Chapter 10. Creating a Table Space	247
Introduction	247
Chapter contents	247
Overview of Chapter	248
Access to the Create Table Space Panel	250
Common Prerequisites for Creating a Table Space	251
Specific Prerequisites and Tasks for a Segmented Table Space	253
Specific Prerequisites and Tasks for a Partitioned Table Space	254
Beginning to Create the Table Space	256
Using Values from an Existing Table Space	257
Using the Default Values for the New Table Space	260
Completing a Non-Partitioned, DB2-Managed Dataset	263
Chapter 11. Creating a Table	265
Introduction	265
Chapter contents	265
Overview of Chapter	266
Access to the Create Table Panel	267
Prerequisites for Creating a Table	268
Beginning to Create the Table	269
Providing Required Values for the New Table	270
Background about Functions and Selects Available	272
Background about Using the Editor	274
Example of How to Use the Editor	276
Defining the Columns for the New Table	278
Adding Comments and Labels to Columns	280
Completing Values for the Tables	281
Creating Synonyms for the Table	282
Chapter 12. Creating an Index and Index Space	283
Introduction	283

Chapter contents	283
Overview of Chapter	284
Access to the Create Index Panel	286
Common Prerequisites for Creating an Index and Index Space	
Space	287
Specific Tasks for Unique or Clustering Indexes	290
Specific Tasks for Partitioned Index and Index Space	291
Beginning to Create the Index	292
Using Values from an Existing Index	294
Using the Default Values for the New Index	296
Selecting Table Columns to Be Indexed	298
Completing a Non-Partitioned, DB2-Managed Dataset	302
Chapter 13. Creating a Storage Group	305
Introduction	305
Chapter contents	305
Overview of Chapter	306
Access to the Create Storage Group Panel	307
Prerequisites for Creating a Storage Group	308
Beginning to Create the Storage Group	309
Using Values and VOLSERS from an Existing Storage Group	310
Modifying Values from an Existing Storage Group	312
Chapter 14. Completing Requirements for VSAM Datasets	315
Introduction	315
Chapter contents	315
Overview of Chapter	316
Partitions and Setting Up the VSAM Dataset	320
Setting Up a DB2-Created and Managed Dataset	322
Setting Up an Existing Dataset That You Manage	325
Defining and Setting Up a VSAM Dataset	328
Background for a Partitioned Index Space	333
Completing Different Values for Partitions	335
Running Utilities	337
Chapter 15. Introduction to DB2 Utilities and Options	339
Introduction	339
Chapter contents	339
Resources for Learning about Topics in This Section	340
Overview of the Chapter	342
Process for Using Utilities and Profiles	343
Accessing Object List Panels to Use with Utilities	344
Accessing RI Data to Use with Utilities	346
Introduction to the Utilities Menus	348
Introduction to the List Option	349
Introduction to Utility Profiles and !DB/WORKBENCH	350
Setting the Active Utility Profile	351

Accessing the Utilities	354
Completing Values for the Utility	356
Processing the Statements Generated	359
Accumulating Statements from Multiple Utilities	361
Chapter 16. DB2 Utilities and Options Available with	
!DB/WORKBENCH	363
Introduction	363
Chapter contents	363
Overview of the Chapter	364
Online Help and Specific Values for Utilities	366
CHECK DATA Utility	367
CHECK INDEX Utility	369
COPY Utility	371
DIAG ABENDS Utility	375
DIAG SYSUTIL Utility	377
DISASTER REC Utility	379
ICCAT Utility	382
MERGE COPY Utility	384
MODIFY Utility	387
QUIESCE Utility	389
REBIND Command	391
REBIND PKGS Command	394
REBIND PLANS Command	397
RECOVER INDEX Utility	400
RECOVER Utility for Table Spaces	402
REORG INDEX Utility	407
REORG Utility for Table Spaces	409
REPORT Utility	411
RUNSTATS Utility	413
STOSPACE Utility	417
Chapter 17. Panels and Fields Used with Multiple Utilities	419
Introduction	419
Chapter contents	419
Overview of the Chapter	420
Online Help and This Chapter	423
Partitioned Objects and Multiple Objects and Utilities	424
Options for DB2 Commands with Utilities	426
General Options for Job Control	428
Specifying Values for the Job Card and Job Statements	429
Values for Identifying and Restarting a Utility	431
Values for the Work Data Sets	432
Values for Sort Work Data Sets	435
Values for Copy Data Sets (LOCAL)	438
Values for Copy Data Sets (RECOVERY)	440
Values for Partitions	442
Modifying the Values for Device (Esoteric) Names	445

How !DB/WORKBENCH Determines Data Set Sizes	446
Chapter 18. Creating or Modifying Profiles for DB2 Utilities and Options	449
Introduction	449
Chapter contents	449
Overview of the Chapter	450
Background about the Profile Option on Utilities Menus	452
Accessing the Utility Profiles Menu	454
Beginning the Process of Creating a Profile	456
Background about Options and Panels in the Profile	458
Providing Values on Panels	460
Saving a Profile	462
Chapter 19. Using User Variables in Utility Profiles	463
Introduction	463
Chapter contents	463
Organization of the Chapter	464
Introduction to Creating User Variables	466
Overview of the Process of Creating and Using User Variables	467
Prerequisites for Creating User Variables	468
Background about the Template for the Panel for User Variables	470
Background about Variables in the Template	472
Characteristics of the Example of Modifying the Template	474
Example of Modifying the Template	476
Introduction to Creating Help for User Variables	478
Prerequisites for Creating Help for User Variables	479
Background about Templates for Help for User Variables	480
Characteristics of the Example of Creating Help for User Variables	482
Example of Creating Help for User Variables	484
Completing the Process of the Example	486
Running Batch and Using JCL	487
Chapter 20. Introduction to !DB/WORKBENCH JCL and Batch	489
Introduction	489
Chapter contents	489
Resources for Learning Topics in This Section	490
Overview of the Chapter	492
Introduction to the Process of Using Batch	493
Prerequisites for Using Batch Processing	494
Reports Produced by the Batch Utility	495
Determining What to Do Next	497

Chapter 21. Requirements for Creating Batch Jobs	499
Introduction	499
Chapter contents	499
Overview of the Chapter	500
Methods to Access Object Lists in Batch	502
Methods to Use Commands in Batch	503
Operation of the END Command in Batch	504
Method to Use the PRNT Command to See Reports of Data	506
Methods to Filter Objects in Batch	507
Availability of General Online Functions in Batch	510
Requirements for Using Commands in Batch Jobs	512
Example of a Simple Batch Job	513
Example of a Simple Batch Job for Utilities	514
Example of a Job Stream with Multiple Batch Jobs	515
Background about Commands to Defer Execution	516
Example of a Simple Job for Deferred Execution	518
Example of Multiple Jobs for Deferred Execution	519
Determining What to Do Next	520
Chapter 22. JCL for Utilities, Batch Jobs, and Extracts	521
Introduction	521
Chapter contents	521
Overview of the Chapter	522
Background about Skeletons for Online and Batch	524
How to Customize a Skeleton or the KTWBATCH Procedure	526
Customization of the KTWUTJCL Skeleton	528
Introduction to !DB/WORKBENCH and the Batch Utility	530
Background about Customization of the KTWBATCH Procedure	531
Making Recommended Modifications for PDS Names	532
Making Recommended Modifications for DD Statements	533
Making Recommended Changes to the Extract ID	534
Making Optional Modifications to Batch Jobs	535
Modifying DD Statements for Batch Extracts	537
Submitting a Batch Job for Execution	538
Chapter 23. Command Dictionary for Batch	539
Introduction	539
Chapter contents	539
Overview of the Chapter	540
General Commands for Specifying Objects in Batch	543
General Commands for Specifying Authorizations in Batch	544
AMS Command and Parameters for Batch	545
BIND Command and Parameters for Batch	547
DISP Command and Parameters for Batch	548
DUMP Command and the Batch Utility	549
END Command and Options in Batch	552
EXEC Command	554

FLIST Command	555
NOEXEC Command and the Batch Utility	558
OUTC Command and Parameters in Batch	559
PDSCOPY Command and the Batch Utilities	562
PDSEEDIT Command and the Batch Utility	564
PDSEXEC Command and the Batch Utility	566
PDSLIST Command and the Batch Utility	567
PRNT Command and Operation in Batch	569
START Command and Parameters in Batch	571
UTIL Command and Functions Supported in Batch	572
Appendix A. Candle Customer Support	575
Customer Support Phone Numbers	577
Index	579

Preface

The *!DB/WORKBENCH Advanced User's Guide* gives you information and direction on how to use SQL and DB2 commands, manage your DB2 system, create objects, run utilities, run batch operations, and use JCL.

The individuals who can benefit from this guide include

- database administrator
- storage administrator
- DB2 systems programmer
- security administrator
- application programmer

How this guide is organized

This guide is organized into five main sections.

- The “*Using SQL and DB2*” section is comprised of two chapters that
 - Assist you in using the dynamic SQL facility in !DB@/WORKBENCH for DB2
 - Show you how to use DB2 commands in !DB/WORKBENCH
- The “*Managing the System*” section contains five chapters that:
 - Explain the use of Authorizations to manage the DB2 system
 - Show you how to estimate DASD usage
 - Guide you in using the DSNZPARMS facility
 - Identify the Housekeeping options in !DB/WORKBENCH
 - Allow you to manage statistics in the STATDBAS table space

- The “*Creating Objects*” section includes seven chapters that:
 - Introduce you to the concept of creating objects
 - Provide information on creating a database
 - Provide information on creating a table space
 - Provide information creating a table
 - Provide information on creating an index and an index space
 - Provide information on creating a storage group
 - Show you how to complete requirements for VSAM datasets
- The “*Running Utilities*” section contains five chapters that
 - Introduce you to DB2 utilities and options
 - Introduce you to DB2 utilities and options available with !DB/WORKBENCH
 - Show you the panels and fields used with multiple utilities
 - Explain how you create or modify profiles for DB2 utilities and options
 - Explain how you use user variables in utility profiles
- The “*Running Batch and Using JCL*” section is comprised of four chapters that
 - Introduce you to !DB/WORKBENCH JCL and batch operations
 - Identify the criteria you need for creating batch operations
 - Explains the JCL you need for utilities and batch operations
 - Lists the command dictionary for batch operations
- The remaining sections contain the Customer Support information and an index. (A complete glossary is provided with the *!DB/WORKBENCH Reference Manual*.)

Conventions

Introduction

Certain conventions have been followed to add to the usability of this document. Examples are given for each convention.

Conventions used in this guide

This guide uses the following conventions.

- | | |
|---------------------------|---|
| Commands | All command names are in uppercase. This includes commands for TSO, DB2, and !DB/WORKBENCH, such as the ALIAS command. |
| Percent (%) sign | A percent (%) sign is shown in front of all CLISTS mentioned in the written instructions for optimum performance. However, the use of the % is optional. For example, the CLIST KTC is preceded by the %, such as %KTC. |
| Function keys | This guide does not refer to specific function keys, such as F10. Because these function keys are user-defined, the guide instructs you to enter the command on the command line.
Example: Type RETURN on the command line and press Enter, or press the appropriate function key. |
| Keyed data | When you are given instructions to key data, the data to be keyed will be in bold, such as RETURN . |
| Panels and figures | The panels and figures in this document are representations. Actual product panels may differ. |
| Revision bars | Revision bars (l) appear on the left margin to identify new or updated material. |
| Variables | If the data is a variable, the data will be in an italic font, such as <i>databasename</i> . The variable name will not necessarily be the correct number of characters allowed by the system, but instead be a descriptive name. |

Adobe Portable Document Format

Introduction

Candle supplies documentation in the Adobe Portable Document Format (PDF). The Adobe Acrobat Reader prints PDF documents with the fonts, formatting, and graphics in the original document. To print a Candle document, do the following:

1. Specify the print options for your system. From the Acrobat Reader Menu bar, select **File > Print Setup...** and make your selections. A setting of 300 dpi is highly recommended as is duplex printing if your printer supports it.
2. To start printing, select **File > Print** on the Acrobat Reader Menu bar.
3. On the Print popup, select one of the **Print Range** options for
 - a single page
 - a range of pages
 - all of the document
4. (Optional) To fit oversize pages to the paper size currently loaded on your printer, select the **Shrink to Fit** option.

Printing problems?

Your printer ultimately determines the print quality of your output. Sometimes printing problems can occur. If you experience printing problems, potential areas to check are:

- settings for your printer and printer driver. (The dpi settings for both your driver and printer should be the same. A setting of 300 dpi is recommended.)
- the printer driver you are using. (You may need a different printer driver or the Universal Printer driver from Adobe. This free printer driver is available at www.adobe.com.)
- the halftone/graphics color adjustment for printing color on black and white printers. (Check the printer properties under **Start > Settings > Printer**. For more information, see the online help for the Acrobat Reader.)
- the amount of available memory in your printer. (Insufficient memory can cause a document or graphics to fail to print.)

For additional information on printing problems, refer to the documentation for your printer or contact your printer manufacturer.

Documentation Set

Introduction

Candle provides a complete set of documentation for !DB/WORKBENCH. Each manual in this documentation set contains a specific type of information to help you use the product.

Candle welcomes your comments and suggestions for changes or additions to the documentation set. A user comment form, located at the back of each manual, provides simple instructions for communicating with Candle's Information Development department. You can also send email to **UserDoc@candle.com**. Please include the product name, version, and book title in the subject line. To order additional manuals, contact Candle Customer Support.

The documentation listed in the following table is available for !DB/WORKBENCH.

Document Number	Document Name	Description
TI51-5840	<i>!DB/Tools Installation and Customization Guide</i>	Gives the requirements and instructions to install all the !DB@/Tools for DB2.
TW54-5847	<i>!DB/WORKBENCH Basic User's Guide</i>	Contains information on <ul style="list-style-type: none">● accessing, using, and exiting !DB/WORKBENCH● online Help● Primary Menu and object list panels● selection lists from object list panels● ISPF facilities within this product● displaying data and filtering object lists● operations on data● !DB/WORKBENCH macros
TW54-5848	<i>!DB/WORKBENCH Advanced User's Guide</i>	Contains information on <ul style="list-style-type: none">● SQL and DB2 commands● managing the system● creating objects● running utilities● running batch operations● JCL.

Document Number	Document Name	Description
TW53-5849	<i>!DB/WORKBENCH Reference Manual</i>	Contains information on <ul style="list-style-type: none"> ● commands and selects in online and batch ● methods for filtering and sorting ● accessing data from the DB2 catalog ● accessing data from the extract ● problem resolution ● abend codes and error messages

Online documentation for BookManager

In addition to the printed versions, !DB/WORKBENCH documentation is available in BookManager format. Subsequent maintenance releases of the product between major releases will contain updated BookManager documentation. See the *!DB/Tools Installation and Customization Guide* for information regarding the files.

For more information about using BookManager, see the appropriate IBM documentation as follows:

IF you want to ...	THEN see ...
put our book on a bookshelf,	<i>IBM BookManager BUILD/MVS Preparing Online Books SC38-2036</i>
display our book,	<i>IBM BookManager READ/MVS Displaying Online Books SC38-2034</i>

Version 500 Changes

Overview

This section describes what is different between !DB/WORKBENCH and previous versions of !DB/WORKBENCH.

Differences in this release

| Primarily, the Version 500 release of !DB/WORKBENCH provides full
| compatibility for DB2 Version 6. However, some DB2 enhancements are not
| exploited by !DB/WORKBENCH Version 500.

DB2 Version 6 contains enhancements for

- defining and manipulating data objects
- conducting e-business
- improving performance and availability of database applications
- managing the database environment
- increasing database and query capacity

This release of !DB/WORKBENCH also continues to support fully earlier releases of DB2 and is also Y2K compliant.

Using SQL and DB2

Chapter 1. Using the Dynamic SQL Facility in !DB/WORKBENCH

Introduction

This chapter provides information about accessing data with SQL from !DB/WORKBENCH. The chapter contains information on executing statements entered directly on the SQL Input panel, executing retrieved statements, and executing statements stored in a PDS member. It also includes information on testing SQL statements.

Chapter contents

Resources for Learning about Features Available from !DB/WORKBENCH	24
Overview of the Chapter	25
Accessing the SQL Input Panel	27
Overview of the SQL Input Panel	28
Using Uppercase and Lowercase	30
Guidelines for SQL Statements Entered Directly	31
Executing Statements Entered Directly	32
Executing Retrieved Statements	33
Recovery from Errors	34
Keyboard Mapping and SQL Statements Stored in a PDS Member	35
Guidelines for SQL Statements Stored in a PDS Member ..	36
Executing Statements Stored in a PDS Member	38
Interpreting the Report Displayed by the System	40
Testing SQL Statements within a Program	42

Resources for Learning about Features Available from !DB/WORKBENCH

Overview

This unit lists resources available for learning about features available from !DB/WORKBENCH.

Resources in this guide

This section contains the options for dynamically entering SQL statements and for using DB2 commands from !DB/WORKBENCH.

Information You Want	Resource
An introduction to the methods you can use to dynamically enter SQL statements within !DB/WORKBENCH	This chapter
An overview of the DB2 commands available with !DB/WORKBENCH and information about authority required and values you must provide to use the DB2 commands	“Using DB2 Commands in !DB/WORKBENCH” on page 45

Overview of the Chapter

Overview

This unit contains introductory information on using !DB/WORKBENCH to interactively input and execute SQL statements. It also contains information on selecting information to match your needs.

Background

In !DB/WORKBENCH, you can interactively input and execute SQL statements. Within !DB/WORKBENCH you can:

- Input and execute statements directly on the SQL Input panel
 - Retrieve one of the nine statements listed on the SQL Input panel (!DB/WORKBENCH saves your last nine successful SQL statements.)
 - Execute statements stored in members of the user PDS or of another PDS
 - Test SQL statements in application programs
-

Organization of information in this chapter

This chapter provides introductory information on the resources for learning about topics in this section.

Organization of the chapter and your needs

Review the chart to select information appropriate for the task you want to perform. Then locate the page number of the units of information that you need for your task.

	Accessing	Panels	Case	Directly	Statements	Retrieved Statements	Errors	Key Board	Guidelines	Stored Statements	Report	Testing
	p.27	p.28	p.30	p.31	p.32	p.33	p.34	p.35	p.36	p.38	p.40	p.42
Access and Exiting	√											
Review background	√	√										
Create and execute statements in the input panel	√	√	√	√	√		√				√	
Retrieve and execute statements directly in the input panel	√	√	√	√		√	√				√	
Create and execute statements saved to a PDS member	√	√	√					√	√	√	√	
Test SQL statements in a program	√	√	√						√			√

Accessing the SQL Input Panel

Overview

This unit contains information about accessing the SQL Input panel from the !DB/WORKBENCH product panels and from the !DB/WORKBENCH Primary Menu. This unit also contains information about exiting the SQL Input panel.

Procedure for accessing the SQL Input panel

Follow these steps to access the SQL Input panel from the !DB/WORKBENCH Primary Menu or from an object list panel.

Step	Action
1	Perform the appropriate action: <ul style="list-style-type: none"> ● On the !DB/WORKBENCH Primary Menu, type SQL in the OPTION field. ● On any object list panel, type SQL on the command line.
2	Press Enter. Result: The system displays a panel that allows you to type SQL statements and execute them.

Procedure for exiting

Follow these steps to exit the SQL Input panel.

Step	Action
1	Type END on the command line.
2	Press Enter. Result: The result varies based on the method of access. <ul style="list-style-type: none"> ● If you accessed the SQL Input panel from the !DB/WORKBENCH Primary Menu, the system returns you to the primary menu. ● If you accessed the SQL Input panel from any other !DB/WORKBENCH panel, the system returns you to that panel.

Overview of the SQL Input Panel

Overview

This unit contains information about the SQL Input panel.

Elements of the SQL Input panel

The illustration shows frequently used elements of the SQL Input panel.

```
Cmd ==> 1
Enter SQL statement(s) below then press ENTER. Separator = ";". END to exit.
Max rows ==> 50 2 Output- 0 3 DB2 ID ==> D23A 4 Plan ==> TW250PL3 5
-CAPS ON 6----- S Q L   S T A T E M E N T S -----
7

--- Select entry from list below by entering its number.-----
8 1. SELECT DBNAME,ICTYPE,SHRLEVEL,TSNAME,DSNUM,START_RBA,ICDATE,ICTIME,DEVTYP
    2.
    3.
    4.
    5.
    6.
    7.
    8.
    9.
```

Elements of the SQL Input panel (continued)

The lists defines the frequently used elements of the SQL Input panel.

- 1** Location for entry of ISPF and !DB/WORKBENCH commands
- 2** Maximum number of rows you want the system to display as a result of an SQL statement (You can change it from its initial value of 50 rows to any value within the program limit. The program limit value varies with the size of the rows.)
- 3** Number of lines displayed for the SQL statement you previously executed
- 4** DB2 system ID (You can change this value if you want to access a different DB2 system.)
- 5** Name of the runtime plan for the statements that you will execute (The system initializes the value you enter to the plan established by your !DB/WORKBENCH installation.)
- 6** Area of display of the CAPS ON or CAPS OFF setting (if set)
- 7** Area for you to begin entering SQL statements (or to enter the exclamation point required for some procedures)
- 8** Beginning text of the nine statements that executed successfully with the last successful statement located in position number 1. (If an error occurs, the system displays error messages in this area of the panel.)

Using Uppercase and Lowercase

Overview

This unit contains background about using uppercase and lowercase to enter values for SQL processing. This unit also contains information about using the CAPS command.

Background

If you want SQL to process your queries correctly, you must enter values in the appropriate case. If you enter all data in lowercase characters, problems can arise. In particular, you must be careful to enter any WHERE criteria for uppercase values in uppercase.

Commands to use

You can use the CAPS command without an operand to reverse the current status. For example, if the system displays capital letters on the panel and you use the CAPS command, the system turns capitalization off. Review the following chart to determine the command to use.

Task You Want to Perform	Command to Use
Type all characters in uppercase	CAPS ON
Type all characters in lowercase	CAPS OFF

Guidelines for SQL Statements Entered Directly

Overview

This unit contains information on the guidelines to follow when entering SQL statements directly.

Guidelines for entry

Follow these guidelines when you enter SQL statements.

- Type the statements continuously or type the statements one line at a time.
- To separate two statements on the same panel, type a semicolon (;) between statements.

Executing Statements Entered Directly

Overview

This unit contains information on the procedure to use to execute statements entered directly.

Procedure for executing statements entered directly

Follow these steps to execute statements entered directly.

Step	Action
1	In the first column position of the area for you to input statements, type the statement.
2	Press Enter. Result: The system executes the statements.

Executing Retrieved Statements

Overview

This unit contains information on the procedure to use to execute retrieved statements.

Order of statements

When you successfully execute a statement, the system places the statement in the first position and deletes the statement in the ninth position.

Procedure for executing retrieved statements

Follow these steps to execute a retrieved statement.

Step	Action
1	In the first column position of the area for you to input statements, type the number of the statement you want.
2	Press Enter. Result: The system displays the statement in the area where you can enter SQL statements.
3	If you want, edit the statement.
4	Press Enter. Result: The system executes the statement.

Recovery from Errors

Overview

This unit contains information on the methods to use to recover from errors.

Options for recovery from errors

When an error occurs, the system responds with an error message. You have these options.

Action You Can Perform	System Response
Type over the input statement and press Enter.	The system re-executes the statement.
Blank out the input statement and press Enter.	The system redisplay the list of selectable statements.
Type END on the command line and press Enter.	The system exits the SQL Facility and redisplay the panel where you used the SQL command.

Keyboard Mapping and SQL Statements Stored in a PDS Member

Overview

This unit provides information about variations in keyboard mapping that can affect operations on the SQL Input panel. This unit is applicable when you want to execute SQL stored in a PDS member.

The exclamation point (!) and the Dynamic SQL facility

The Dynamic SQL facility supports execution of SQL statements stored in a PDS member. To use this feature, you must type an exclamation point (!) on the SQL Input panel in column 1 of the input area. The hexadecimal code for the exclamation point is X'5A'. If you are using a terminal with, for example, a code page that supports your national language, you will need to determine the character in your code page that corresponds to X'5A' and use that character in place of the exclamation point.

Guidelines for SQL Statements Stored in a PDS Member

Overview

This unit contains information about the guidelines to follow for statements stored in a PDS member.

Errors and statements stored in a PDS member

If the statement results in an error message, the system issues a ROLLBACK.

Syntax for input of statements to a PDS member

Follow these syntax rules for statements stored in a PDS member.

- Input members in SPUFI format.
 - Input a maximum length of 8100 characters (not counting blanks at the beginning and end of a statement unless they are within a quoted string).
 - Use any number of imbedded blanks in the statement.
 - Do not input more than one select statement. (The output is unpredictable.)
-

Additional principles for quoted strings

You can use a quoted string as a comment or input for character fields; however, you must duplicate imbedded quotes. (!DB/WORKBENCH passes the entire string to DB2 exactly as you input it.)

Also, if you use a quoted string that is longer than a single input record, place a non-blank character in Column 72. Then, continue to type the string in the first column of the next line.

Symbols required for statements in the PDS member

Review the following chart for the symbols to use for inputting SQL statements in the member of the PDS.

Task You Want to Perform	Aciton
Use a comment line.	Type — (two dashes) before the lines you do not want the system to read at any position. (Must be in columns 1 and 2.)
End the current statement and begin a comment line.	Type * (asterisk) at the beginning of the comment line. (Must be in column 1.)
End a statement.	Type ; (semicolon) at the end of the statement.

Column positions

Review the following chart to determine the column position.

Column Position	Purpose
1-72	Data that the system will use
72	The continuation column for statements with imbedded quoted strings
73-80	Sequence numbers

Executing Statements Stored in a PDS Member

Overview

This unit contains information about two ways to execute statements stored in a PDS member. You can execute statements in the !DB/WORKBENCH PDS or you can access a panel that allows you to specify a member in a different PDS. This unit contains instructions and background.

Procedure to execute a statement stored in the user PDS

Follow these steps to execute a statement stored in the user PDS.

Step	Action
1	On the SQL Input panel, type ! (or the equivalent symbol) followed by the name of the member of the !DB/WORKBENCH user PDS. Example: If the name of the member is TEST1, then type !TEST1 in the area of input.
2	Press Enter. Result: The system reads and executes all statements.

Background about changing the PDS member

Follow these guidelines when changing the names of PDS.

- The PDS must exist.
- Use standard TSO rules.
- If the dataset name begins with your TSO prefix (not necessarily your USERID), then you do not need to enter the name in quotes.
- If you do not use quotes, then the system prefixes the name with your TSO prefix.
- If the name is in quotes, you must provide the fully qualified name.

Procedure to execute a statement stored in a different PDS

Follow these steps to execute a statement stored in a different PDS.

Step	Action
1	In the first column of the area you want to input statements, type ! on the SQL Input panel.
2	Press Enter. Result: The system displays a panel that allows you to change the PDS name.
3	Type the new PDS name.
4	Perform the appropriate action: <ul style="list-style-type: none"> ● If you know the name of the PDS member, type the name in the member field and press Enter. ● If you do not know the name, press Enter. Result: The result varies with your previous action: <ul style="list-style-type: none"> ● If you typed the member name, the system reads and passes all statements in that member for DB2 to execute. ● If you did not type a member name, the system displays a list of member names (if the PDS is the !DB/WORKBENCH user PDS).
5	If you are selecting from the list of the members in the PDS, type the appropriate value in the select field: <ul style="list-style-type: none"> ● Type S to execute a specific member. ● Type D to delete the member. ● Type E to edit the member. ● Type P to print the member.
6	Press Enter.

Interpreting the Report Displayed by the System

Overview

This unit contains information on interpreting the report displayed by the system.

Interpreting the output

!DB/WORKBENCH follows these conventions in displaying output from the SQL statements you execute.

Item Displayed	Source
Label for the column	Column name or label (if defined)
Field width of the column	The shortest of these values: <ul style="list-style-type: none">• Field length• Length of label or of the name of the column• 36 characters (If a column consists of more than 36 characters, the system displays only 36.)

Displaying the output

Use these methods to display other columns of output that will not fit on the screen.

- Display the rows of data on the panel by scrolling.
- Display columns of data that may not be visible by shifting to the right or left. (The system moves a full page at a time, but it always breaks the display on the boundary of a column and not in the middle of a column.)

Difference between statements executed directly and stored

Output from SQL statements that you have stored varies from statements that you have entered directly in the following two ways:

- Output from stored SQL is displayed only after the input statements have been processed.
- Output after the first 80,000 characters is ignored.

Testing SQL Statements within a Program

Overview

This unit contains information about testing SQL statements you have imbedded in your programs.

Options

You can use any library in the PDS name field except a LOAD library. For example, you can specify any library that contains SQL, such as a CLIST library, a COBOL source library, a PL/1 source library, or a CNTL dataset.

Procedure for accessing the ISPF Editor

Follow these steps to access the ISPF Editor.

Step	Action
1	On the SQL Input panel, type ! in the first column position of the area to input statements.
2	Press Enter. Result: The system displays the panel that allows you to specify the name of the member of the !DB/WORKBENCH user PDS.
3	Type the name of the library you want to access.
4	Press Enter.
5	Type E in the select field for the name of the member you want to edit.
6	Press Enter. Result: The system displays the contents of the member.

Background on specifying a set of lines within the ISPF Editor

Review the chart to specify the statements you want to test in your member.

Segment You Want to Test	Action to Perform
Single lines	Type Q on the line of the statement.
Multiple lines	Type QQ on the first line of the statement. Type QQ on the last line of the statement.

Procedure for editing and testing the statements

Follow these steps to use the ISPF Editor to specify a set of lines for editing and to test the statements.

Step	Action
1	Use the ISPF Editor to specify a set of lines you want the system to execute.
2	Type KTCSQL . Result: The result varies with your action. <ul style="list-style-type: none"> ● If DB2 finds an error in the statement, the system displays a panel with error messages. ● If the statement is correct, the system displays the results.
3	If an error occurs, press Enter to return to the ISPF Editor to edit and execute the statement.

Chapter 2.

Using DB2 Commands in !DB/WORKBENCH

Introduction

This chapter provides instructions, an overview of the options available to use DB2 commands within !DB/WORKBENCH, and prerequisite authority and values needed to use the DB2 commands. Each chapter in this section is independent and can be used without referencing the others.

Chapter contents

Overview of the Chapter	46
Resources for Using DB2 Commands	49
Accessing the DB2 Commands Menu	50
Accessing Options on the DB2 Commands Menu	52
Overview of DB2 Commands and Your Tasks	53
Overview of Operation of the Options and Output	59
Introduction to DB2 Commands for Databases	61
Introduction to the Command to Stop DB2	66
Introduction to DB2 Commands for DDF	67
Introduction to DB2 Commands for RLF	69
Introduction to DB2 Commands for Threads	72
Introduction to DB2 Commands for Traces	77
Introduction to DB2 Commands for Utilities	82
Introduction to DB2 Commands for the Archive Log	83
Introduction to DB2 Commands for Buffer Pools	85
Introduction to DB2 Commands to Recover BSDS	91
Introduction to DB2 Commands for Data Sharing	92

Overview of the Chapter

Background about DB2 commands and !DB/WORKBENCH

!DB/WORKBENCH supports a variety of DB2 commands. You can use these commands from the DB2 Commands menu without leaving !DB/WORKBENCH. Using the commands you can manage databases, the distributed data facility (DDF), the resource limit facility (RLF), DB2 data sharing, and utilities.

You can display data and statistics about threads, and you can recover threads left in doubt after DB2 is restarted. You can also perform the six traces available with DB2 commands, and you can display, modify, start, and stop these traces.

You can also archive the current log. You can stop DB2. You can recover from problems in DB2 from !DB/WORKBENCH. You can use the RECOVER BSDS command to copy the valid bootstrap dataset (BSDS) in a newly-allocated dataset so you can restore the dual mode for BSDS.

Organization of this chapter and your needs

The chapter focuses on the options available to use DB2 commands without leaving !DB/WORKBENCH. It provides general instructions on accessing the DB2 Commands menu and options. It includes an overview of the commands available and a series of units on the authority required and values you will need to provide when you use the option.

Background about DB2 commands and batch processing

This chapter includes information applicable to both online and batch processing with these DB2 commands for the DISPLAY DATABASE and START DATABASE command:

- Display Database Status (DD)
- Start a Database (SD)
- Stop a Database (PD)

Additional information is available for batch processing for the DISPLAY DATABASE and START DATABASE commands. For more information see “Command Dictionary for Batch” on page 539.

Organization of information if you want to perform a task

If you want to perform a specific task using DB2 commands available from !DB/WORKBENCH, review the chart to select information appropriate for your task.

	Resources	Accessing the Menu	Accessing Options	Options Available	Operations/Output	Databases	DB2	DDF	RLF	Threads	Traces	Utilities	Archive Log	Buffer Pools	BSDS	Data Sharing
	p.49	p.50	p.52	p.53	p.59	p.61	p.66	p.67	p.69	p.72	p.77	p.82	p.83	p.85	p.91	P.92
Review resources available for using DB2 commands	√															
Review options for using DB2 commands within !DB/WORKBENCH				√												
Use DB2 commands for managing databases	√	√	√		√	√										
Use the DB2 command to stop DB2	√	√	√		√		√									
Use DB2 commands to manage DDF	√	√	√		√			√								
Use DB2 commands to manage RLF	√	√	√		√				√							
Use DB2 commands to manage threads	√	√	√		√					√						
Use DB2 commands to manage traces	√	√	√		√						√					
Use DB2 commands to manage utilities	√	√	√		√							√				
Use DB2 commands to manage the archive log	√	√	√		√								√			
Use DB2 commands to manage buffer pools	√	√	√		√									√		
Use the DB2 command to recover BSDS	√	√	√		√										√	
Use DB2 commands to manage data sharing	√	√	√		√											√

Organization of information if you want to use a specific DB2 command

If you want to use a specific DB2 command from !DB/WORKBENCH, review the chart for information appropriate for the command.

	Resources	Accessing the Menu	Accessing Options	Options Available	Operations/Output	Databases	DB2	DDF	RLF	Threads	Traces	Utilities	Archive Log	Buffer Pools	BSDS	Data Sharing
	p.49	p.50	p.52	p.53	p.59	p.61	p.66	p.67	p.69	p.72	p.77	p.82	p.83	p.85	p.91	p.92
ALTER BUFFERPOOL	√	√	√		√									√		
ALTER GROUP BUFFERPOOL	√	√	√		√									√		
ARCHIVE LOG	√	√	√		√								√			
CANCEL DDF THREAD	√	√	√		√					√						
DISPLAY ARCHIVE	√	√	√		√								√			
DISPLAY BUFFERPOOL	√	√	√		√									√		
DISPLAY DATABASE	√	√	√		√	√										
DISPLAY GROUP	√	√	√		√											√
DISPLAY GROUP BUFFERPOOL	√	√	√		√									√		
DISPLAY LOCATION	√	√	√		√											√
DISPLAY RLIMIT	√	√	√		√				√							
DISPLAY THREAD	√	√	√		√					√						
DISPLAY TRACE	√	√	√		√						√					
DISPLAY UTILITY	√	√	√		√							√				
MODIFY TRACE	√	√	√		√						√					
RECOVER BSDS	√	√	√		√										√	
RECOVER INDOUBT	√	√	√		√					√						
RESET GENERIC LU	√	√	√		√											√
RESET INDOUBT	√	√	√		√					√						
SET ARCHIVE	√	√	√		√								√			
START DATABASE	√	√	√		√	√										
START DDF	√	√	√		√			√								
START RLIMIT	√	√	√		√				√							
START TRACE	√	√	√		√						√					
STOP DATABASE	√	√	√		√	√										
STOP DB2	√	√	√		√		√									
STOP DDF	√	√	√		√			√								
STOP RLIMIT	√	√	√		√				√							
STOP TRACE	√	√	√		√						√					
TERM UTILITY	√	√	√		√							√				

Resources for Using DB2 Commands

Overview

This unit provides resources for using DB2 commands. It includes information about the resources in !DB/WORKBENCH and in the !DB/WORKBENCH online Help.

!DB/WORKBENCH panels for entry

When a DB2 command you want to use requires that you provide values in fields, !DB/WORKBENCH provides brief instructions on the panel and defines the fields.

Online Help and entry of values

Online Help provides a greater level of detail about the values you enter. This Help panel for STOP TRACE–GLOBAL is typical.

```

Cmd ==>
----- HELP FOR DISPLAY OR STOP TRACE - GLOBAL -----
Press END to return, ENTER to continue
or TAB to highlighted text and press type to display more information.

Authids          To limit the display or stop traces started for particular
                  authorization identifiers, type up to 8 IDs separated by
                  commas. If you typed more than 1 ID, you can only type 1
                  value for Plans. Default is * for all IDs.

Classes          To limit the display or stop traces for particular
                  classes, type the class values separated by commas.
                  Default is 1, 2, or 3.

RMIDs            To limit the display or stop traces for particular resource
                  managers, type up to 8 manager identifiers or type ? to
                  display RESOURCE MANAGER ID LIST. Default is * for all
                  managers.

-----
GLOSSARY  INDEX  COMMANDS  FILTERS  WINDOWS  SELECTS  <-PREV  MORE->

```

Accessing the DB2 Commands Menu

Overview

This unit contains information on methods to access the DB2 Commands menu. It includes these methods:

- Accessing from the !DB/WORKBENCH Primary Menu
- Accessing from an object list panel

Prerequisites for accessing from the !DB/WORKBENCH Primary Menu

Before you can use commands on the DB2 Commands menu, you must be using or have previously used one of the object lists. For example, if you display the !DB/WORKBENCH Primary Menu for the first time in a session and you immediately use the – (dash) command, !DB/WORKBENCH displays the DB2 Commands menu. If, however, you try to use one of the options, !DB/WORKBENCH displays an error message.

Accessing the DB2 Commands menu from the !DB/WORKBENCH Primary Menu

Follow these steps to access the DB2 Commands menu from the !DB/WORKBENCH Primary Menu. (Be sure you have previously used one or more object lists before accessing the command from the !DB/WORKBENCH Primary Menu.)

Step	Action
1	On the !DB/WORKBENCH Primary Menu, type – (dash) in the Option field.
2	Press Enter. Result: The system displays a menu listing DB2 commands available in !DB/WORKBENCH.

Accessing the DB2 Commands menu from an object list panel

Follow these steps to access the DB2 Commands menu from an object list panel.

Step	Action
1	On an object list panel, type – (dash) on the command line.
2	Press Enter. Result: The system displays a menu listing DB2 commands available in !DB/WORKBENCH.

Accessing Options on the DB2 Commands Menu

Overview

This unit contains information on accessing options from the DB2 Commands menu.

Accessing an option from the DB2 Commands menu

Follow these steps to access an option from the DB2 Commands menu.

Step	Action
1	On the DB2 Commands menu, type the two or four character abbreviation of the DB2 command on the option line. Example: If you want to display databases (DD), type DD in the Function field.
2	Press Enter. Result: The system displays the panel for the DB2 command you select.

Overview of DB2 Commands and Your Tasks

Overview

This unit shows the menu available with DB2 commands and surveys DB2 commands available for managing databases, the stopping of DB2, the DDF facility, the RLF facility, threads, traces, utilities, the archive log, buffer pools, the BSDS (bootstrap dataset), and data sharing.

Background about functions available from the menu

When you access the menu for DB2 commands, you have these functions available.

```

DB2 COMMANDS SELECT ----- DB/WORKBENCH --- DB2=DB51-----
Cmd ==>                                     99/12/18 09:41

Option ==>

Archive          DB2          Procedures      Trace
AR LOG*         P2 STOP*       DP DISPLAY*    DT DISPLAY*
DA DISPLAY*     Database      SP START*      MT MODIFY*
SA SET*         DD DISPLAY    PP STOP*       ST START*
                SD START
Bufferpool      RLFacility
AB ALTER*      PD STOP       DR DISPLAY     Utility
AGBP ALT GRP*  SR START     SR START       DU DISPLAY
DB DISPLAY*    DDFacility   PR STOP       PU STOP
DGBP DISP GRP* SF START*
                PF STOP*
Thread
BSDS           CT CANCEL
RB RECOVER *   DC DISP Connection
                RI RECOVER Indoubt*
                RS RESET Indoubt*

                Data Sharing
                DG DISPLAY GROUP*
                DL DISP Location*
                RG RESTET GENERICLU*

* ( SYSOPR/SYSCTRL/SYSADM or appropriate authorization is needed )
    
```

Options available for managing databases

This chart lists the options available for managing databases.

Task You Want to Perform	Option to Select
Display the status of the database (Use the DISPLAY DATABASE command.)	DD–Display Database Status
Start a list of databases or table spaces and indexes (Use the START DATABASE command.)	SD–Start a Database
Stop a list of databases or table spaces and indexes (Use the STOP DATABASE command.)	PD–Stop a Database

Options available for managing DB2

This chart lists the options available for managing DB2.

Task You Want to Perform	Option to Select
Stop the DB2 subsystem (Use the STOP DB2 command.)	P2–Stop DB2

Options for managing the distributed data facility

This chart lists the options available for managing the DDF (distributed data facility).

Task You Want to Perform	Option to Select
Start the distributed data facility (Use the START DDF command.)	SF–START DDF
Stop the distributed data facility (Use the STOP DDF command.)	PF–STOP DDF

Options for managing the resource limit facility

This chart lists the options available for managing the resource limit facility (RLF).

Task You Want to Perform	Option to Select
Display the status of the resource limit facility (Use the DISPLAY RLIMIT command.)	DR–Display RLIMIT
Start the resource limit facility (Use the START RLIMIT command.)	SR–Start RLIMIT
Stop the resource limit facility (Use the STOP RLIMIT command.)	PR–Stop RLIMIT

Options for managing threads

This chart lists the options available for managing threads.

Task You Want to Perform	Option to Select
Cancel processing for threads (Use the CANCEL DDF THREAD command.)	CT–Cancel DDF Thread
Display information about DB2 and distributed subsystem connections (Use the DISPLAY THREAD command.)	DC–Display Connections (Thread)
Recover threads left in doubt after DB2 is restarted (Use the RECOVER INDOUBT command.)	RI–Recover Indoubt
Purge information for in doubt threads in the report generated by the DISPLAY THREAD command (Use the RESET INDOUBT command.)	RS–Reset Indoubt

Options for managing traces

This chart lists the options available for managing traces.

Task You Want to Perform	Option to Select
Display the status of DB2 traces (Use the DISPLAY TRACE command.)	DT–Display Active Traces
Modify the trace events being traced for a specific active trace (Use the MODIFY TRACE command.)	MT–Modify Trace
Start the DB2 trace (Use the START TRACE command.)	ST–Start a Trace
Stop the DB2 trace (Use the STOP TRACE command.)	PT–Stop a Trace

Options for managing DB2 utilities

This chart lists the options available for managing DB2 utilities.

Task You Want to Perform	Option to Select
Display the status of a DB2 utility (Use the DISPLAY UTILITY command.)	DU–Display Utility Status
Stop a utility that you do not want to restart or that you cannot start (Use the TERM UTILITY command.)	PU–Stop a Utility

Options for managing the Archive Log

This chart lists the options available for managing the Archive Log.

Task You Want to Perform	Option to Select
Archive the current log (Use the ARCHIVE LOG command.)	AR–Archive Log
Display information about processing of the archive log (Use the DISPLAY ARCHIVE command.)	DA–Display Archive
Set the limit for the deallocation time and number of tape units for the archive log (Use the SET ARCHIVE command.)	SA–Set Archive

Options for managing buffer pools

This chart lists the options available for managing buffer pools.

Task You Want to Perform	Option to Select
Alter attributes for active or inactive buffer pools (Use the ALTER BUFFERPOOL command.)	AB–Alter Bufferpool
Alter attributes for active or inactive group buffer pools (Use the ALTER GROUP BUFFERPOOL command.)	AGBP–Alter Group Bufferpool
Display information about one or more active or inactive buffer pools (Use the DISPLAY BUFFERPOOL command.)	DB–Display Bufferpool
Display information about one or more active or inactive group buffer pools (Use the DISPLAY GROUP BUFFERPOOL command.)	DGBP–Display Group Bufferpool

Options for managing BSDS recovery

This chart lists the options available for managing bootstrap dataset (BSDS) recovery.

Task You Want to Perform	Option to Select
Make a copy of the valid bootstrap data set in a newly allocated dataset and therefore restore the dual mode of the BSDS. (Use the RECOVER BSDS command.)	RB–Recover BSDS

Options for managing data sharing

This chart lists the options available for managing DB2 data sharing.

Task You Want to Perform	Option to Select
Display DB2 data sharing group information (Use the DISPLAY GROUP command.)	DG–Display Group
Display statistics about threads and connections between remote DB2 subsystems (Use the DISPLAY LOCATIONS command.)	DL–Display Locations
Purge VTAM information from the coupling facility regarding partners of this DB2 subsystem. (Use the RESET GENERIC LU command.)	RG–Reset Generic LU

Overview of Operation of the Options and Output

Overview

This unit provides information on these topics about the options available on the DB2 Commands menu and the output generated for the options:

- Operating the DB2 Commands menu options
- Providing values for options on a panel
- Interpreting output from the options

Overview of operation of the options

When you access the DB2 Commands menu and select an option, the system varies in operation depending upon the command. If DB2 requires that you specify objects or clauses for the command, !DB/WORKBENCH displays a panel so that you can provide these values. When you provide these values (if any) and press Enter, the system executes the DB2 command and displays the output.

Resources for successfully providing values on a panel

This chapter provides information about the authority required to use the option and values you provide for the option. In this way, you have the choice of determining when you need to use the command without accessing the option.

Online Help for the panel provides specific information about acceptable values for fields. In this way, you have detailed information immediately accessible.

Overview of the output from the options

With !DB/WORKBENCH, the output for the option consists of two parts of the panel separated by a dashed line:

- The top part contains the DB2 command generated when you selected the option and provided values
- The bottom part contains the results from the command

Resources for successfully interpreting output

In addition to resources available with DB2 documentation, you can also use online Help with !DB/WORKBENCH to interpret output from the command. The online Help panel you can display for the DISPLAY TRACES (the Display Active Traces option) command is typical.

```

Cmd ==>
----- HELP FOR DISPLAY TRACE COMMAND OUTPUT -----
Press END to return, ENTER+to continue
or TAB to highlighted text and press ENTER to display more information.

      INPUT COMMAND -
      -DISPLAY TRACE (*          ) +
      DEST( SMF) +
      PLAN(*) +
      AUTHID(*) +
      RMID(*) +
      COMMENT('DI
Trace number | Trace classes | Qualifier
-----> 01 STAT 01,03,04 <--
      .---> ACCTG 01
      DSNW127I ! CURRENT TR   ACTIVITY IS
      TNO TYPE CLASS   DEST QUAL
      .---> ACCTG 01         SMF NO <-----
      .---> ACCTG 01         SMF <-----
      PLAY TRACE S | ARY DATA*****
Trace type:
Global, Performance, Accounting,
Statistical, Audit, or Monitor | h | Output destination:
SMF, GTF, SRV, RES, or OP1-8
-----
GLOSSARY INDEX  COMMANDS  FILTERS  WINDOWS  SELECTS
    
```

Introduction to DB2 Commands for Databases

Overview

This unit contains introductory information on the !DB/WORKBENCH options for the DB2 commands for databases. It also provides information you will need to use the options.

Description of the process of using DB2 commands

With the !DB/WORKBENCH options for the DB2 commands to display, start, and stop databases, you select the option and the system displays a panel listing values to enter.

You can use selects with lists of databases or the associated table spaces and indexes. When you have entered the values and selected the databases or table spaces and indexes, press Enter, and the system executes the command and displays the output and the command statements it generates.

Options for selecting from lists with commands for databases

This chart lists the options available for selecting from lists with commands for databases.

Task You Want to Perform	Action
Select an object to be used with the command	Type S in the Select field for the object
Display the table spaces and index spaces for a database	Type L in the Select field for the database
Remove a database that you previously selected	Type D in the Select field for the database

Options that manage the database

Use these !DB/WORKBENCH options to manage the database:

- Display Database Status (DD)
- Start a Database (SD)
- Stop a Database (PD)

Authority required for Display Database Status

If you want to use the option for Display Database Status, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- DISPLAY privilege
- DISPLAYDB privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database

Values required for Display Database Status

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Database name	"*" or database name (blank for list)
Space name	"*" or space name (blank for list) "–" for database only (no spaces)
Partition No.	Restricts the DB2 command to a single partition (DB2 Ver. 3.1 or higher)
Partition Range	Restricts the DB2 command to a range of partitions (DB2 Ver. 5.1 or higher)
Use	Y or N for usage (correlations & connections)
Locks	Y or N for Lock info (overrides Use)
After	Y or N if the name(s) are start of a list.
Active	Y or N if only want active databases & spaces
Restrict	Y or N if you only want those with restricted use
Claimers	Y or N (V3.1 or higher) Overrides LOCK and USE

Values required for Display Database Status (Continued)

Field Name	Description
LPL	Y or N if you want to list the pages that are logically in error on the logical page list (LPL)
No Limit Output Lines	Y if you want no limit on output Max number of output lines

Authority required for Start a Database

If you want to use the option for Start a Database, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM, DBCTRL, or DBMAINT authority for each database
- STARTDB privilege for each database

Values required for Start a Database

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Database name	"*" or database name (blank for list)
Space name	"*" or space name (blank for list) "-" for no spaces (start database only)
Access	W Read or Write R Read only U Utility use only F Force reset
Partition No.	Restricts the DB2 command to a single partition (DB2 Ver. 3.1 or higher)
Partition Range	Restricts the DB2 command to a range of partitions (DB2 Ver. 5.1 or higher)

Authority required for Stop a Database

If you want to use the option for Stop a Database, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM, DBCTRL, or DBMAINT authority for each database
- STOPDB privilege for each database

Values required for Stop a Database

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Database name	"*" or database name (blank for list)
Space name	"*" or space name (blank for list) "-" for database only (no spaces)
Partition No.	Restricts the DB2 command to a single partition (DB2 Ver. 3.1 or higher)
Partition Range	Restricts the DB2 command to a range of partitions (DB2 Ver. 5.1 or higher)
AT(COMMIT)	Y/N Permits Stop to occur before SQL Applications have completed (DB2 Ver.3.1 or higher)

Introduction to the Command to Stop DB2

Overview

This unit contains introductory information on the !DB/WORKBENCH option for the DB2 command used to stop DB2. It also provides information you will need to use the option.

Authority required for Stop DB2

If you want to use the option for Stop DB2, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
 - STOPALL privilege
-

Values required for Stop DB2

To access the DB2 command available with the option, you provide a value for this field. The fields are displayed in the order they appear on the panel.

Field Name	Description
Quiesce	Q to Quiesce F to Force

Introduction to DB2 Commands for DDF

Overview

This unit contains introductory information on the !DB/WORKBENCH options for the DB2 commands for distributed data facility (DDF). It also provides information you will need to use the options.

Options that manage DDF

Use these !DB/WORKBENCH options to manage DDF:

- Start DDF (SF)
 - Stop DDF (PF)
-

Authority required for Start DDF

If you want to use the option for Start DDF, you must have one of these authorizations:

- SYSADM authority
 - SYSCTRL authority
 - SYSOPR authority
-

Values required for Start DDF

You do not need to provide values to begin to use the Start DDF option.

Authority required for Stop DDF

If you want to use the option for Stop DDF, you must have one of these authorizations:

- SYSADM authority
- SYSCTRL authority
- SYSOPR authority

Values required for Stop DDF

To access the DB2 command available with the option, you provide a value for this field.

Field Name	Description
Mode	STOP DDF Mode F Force Q Quiesce

Introduction to DB2 Commands for RLF

Overview

This unit contains introductory information on the !DB/WORKBENCH options for the DB2 commands for resource limit facility (RLF). It also provides information you will need to use the options.

Options that manage RLF

Use these !DB/WORKBENCH options to manage RLF:

- Display RLIMIT (DR)
 - Start RLIMIT (SR)
 - Stop RLIMIT (PR)
-

Authority required for Display RLIMIT

If you want to use the option for Display RLIMIT, you must have one of these authorizations:

- SYSADM authority
 - SYSCTRL authority
 - SYSOPR authority
-

Values required for Display RLIMIT

You do not need to provide values to begin to use the Display RLIMIT option.

Authority required for Start RLIMIT

If you want to use the option for Start RLIMIT, you must have one of these authorizations:

- SYSADM authority
- SYSCTRL authority
- SYSOPR authority

Values required for Start RLIMIT

To access the DB2 command available with the option, you provide a value for this field.

Field Name	Description
Limit table ID	ID of the resource limit specification table

Authority required for Stop RLIMIT

If you want to use the option for Display RLIMIT, you must have one of these authorizations:

- SYSADM authority
- SYSCTRL authority
- SYSOPR authority

Values required for Stop RLIMIT

You do not need to provide values to begin to use the Stop RLIMIT option.

Introduction to DB2 Commands for Threads

Overview

This unit contains introductory information on the !DB/WORKBENCH options for the DB2 commands for threads. It also provides information you will need to use the options.

Options that manage threads

Use these !DB/WORKBENCH options to manage threads:

- Cancel DDF Thread (CT)
 - Display Connections (Threads) (DC)
 - Recover Indoubt (RI)
 - Reset Indoubt (RS)
-

Background about output of these options

The options that manage threads produce output that is dynamic. Changes from the values displayed may occur in the time required to display the output.

Authority required for Cancel DDF Thread

If you want to use the option for Cancel DDF Thread, you must have one of these authorizations:

- SYSADM authority
- SYSCTRL authority
- SYSOPR authority

Values required for Cancel DDF Thread

To access the DB2 command available with the option, you provide values for these fields. The fields are in the order they appear on the panel.

Field Name	Description
Token	1–5 digit number
Network ID	(1–8 Chars.)
LUW Instance No.	(12 Hex chars.)
Commit Seq. No.	(4 Hex chars.)

Authority required for Display Connections

If you want to use the option for Display Connections, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- DISPLAY privilege

Values required for Display Connections

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Connection ID	"*" or thread name
Type	Type of thread to display A Active I Indoubt B Both
Location	Location: (DDF only, "*" or Location Name) "*" or name
Detail?	Y/N
Token	1-5 digit number
Network ID	(1-8 Chars.)
LUW Instance No.	(12 Hex chars.)
Commit Seq. No.	(4 Hex chars.)

Authority required for Recover Indoubt

If you want to use the option for Recover Indoubt, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- RECOVER privilege

Values required for Recover Indoubt

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Connection name	Recovers all threads for the connections
Action	COMMIT or ABORT
Correlation ID or Network ID	For specific thread, type thread ID For all threads, type * If correlation IDs are not unique, type network ID for thread
Token	1–digit number
Network ID for LUWID	(1–8 Chars.)
LUW Instance number	(12 Hex chars)
Commit Seq. No.	(4 Hex chars.)

Authority required for Reset Indoubt

If you want to use the option for Reset Indoubt, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- RECOVER privilege

Values required for Reset Indoubt

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
LUnicode	1–8 character name or * for all qualifying locations
Force	Y Simulates COLD Start
Location	DB2 DBMS or non-DB2 Server
Token	1–5 digit number
Network ID	(1–8 Chars.)
LUW Instance No.	(12 Hex chars.)
Commit Seq. No.	(4 Hex chars.)

Introduction to DB2 Commands for Traces

Overview

This unit contains introductory information on the !DB/WORKBENCH options for the DB2 commands for traces. It also provides information you will need to use the options.

Because the values required for three of these options are numerous and vary with the option, the unit includes both a list and description of values and charts showing which values are required for different options.

Options that manage traces

Use these !DB/WORKBENCH options to manage traces:

- Display Active Trace (DT)
 - Modify Trace (MT)
 - Start a Trace (ST)
 - Stop a Trace (PT)
-

Description of the process of using DB2 commands for traces

With the !DB/WORKBENCH options for DB2 commands for managing traces, you select the option and the system displays a menu listing the six types of data that DB2 traces can record. The Display and Stop Trace options also list menu options to display all traces.

When you select the type of trace from the menu, the system displays a panel for you to provide values in fields so !DB/WORKBENCH can generate the DB2 commands statement for each trace. (The display and stop options provide an additional panel which you can use if you want to provide additional restrictions for your trace. If you do not want to restrict the trace further, press Enter.) When you complete the values in the fields, press Enter and the system executes the command and displays the command statements it generated and the output.

Options for selecting from lists of objects

You can select from a list of RMIDs by using a question mark (?) in the field. The system displays a list. You can then type an S in the Select field for each one you want.

Options on the menus for traces

The menus for displays, modifications, starting, and stopping traces provide traces for the same type of data. The chart lists the option letter, the clause for the data, and a brief description from the menu.

Option	Clause	Description
G	GLOBAL	DB2 Servicing - traces from entire system
P	PERFM	Performance tuning - specific events
A	ACCTG	Accounting - records for each thread
S	STAT	Statistical data from DB2 components
U	AUDIT	Audit data from DB2 components
M	MONITOR	Collect monitor data (for starting traces) Monitor data from DB2 components
*	ALL	Display all traces (for Display and Stop options only)

Authority required for all types of traces

If you want to use the options for traces, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- TRACE privilege

Options for comments

The panels provide an optional comment field for the trace. The field length is a maximum of 60 characters.

Description of values for displaying, monitoring, starting, and stopping traces

The chart provides field names and descriptions for values used with the options for displaying, monitoring, starting, and stopping traces. (Charts listing which fields are required for which options are included on subsequent pages in this unit. The fields on the charts are in the same order as listed here.)

Field Name	Description
SMF	Destination (MVS System Management Facility)
GTF	Destination (Generalized Trace Facility)
SRV	Destination (serviceability routine reserved for problem diagnosis)
RES	Destination (A resident table in the Common Service Area)
OPX	Destination (the next available buffer)
OPn	Destination (A location belonging to a specific application program)
Plans	1–8 names, separated by commas (* = default)
Authids	1–8 names, separated by commas (* = default) Max: 8 Plans and 1 Authid or 1 Plan & 8 Authids
Classes	Classes: List of IDs (1–5) Default = 1,2,3
RMIDs	RMIDs: Up to 8 IDs (1,2,etc. "?" for list)
Trace IDs	Trace IDs: Up to 8 IDs
Location	Location: (DDF only, separated by commas or *)

Values required for starting traces

The chart lists the values required on panels if you want to start a trace for any of the six possible DB2 traces. The fields are displayed in the order they appear on the panel.

	GLOBAL	PERFM	ACCTG	STAT	AUDIT	MONITOR
SMF	√	√	√	√	√	√
GTF	√	√	√	√	√	√
SRV (not valid for DB2 V2.3 or higher)	√	√	√	√	√	√
OPX	√	√	√	√	√	√
RES	√					
OPn	√	√	√	√	√	√
Plans	√	√	√		√	√
Authids	√	√	√		√	√
Classes	√	√	√	√	√	√
RMIDs	√	√				√
Location	√	√	√		√	√

Values required for monitoring traces

These values are required to use the Monitor option. The fields are displayed in the order they appear on the panel.

Field Name	Description
Classes	Classes: List of IDs (1–5)
Trace IDs	Specify a trace ID

Values required for displaying and stopping traces

The chart lists the values required on panels if you want to display or stop a trace for any of the six possible DB2 traces. The fields are displayed in the order they appear on the panel.

	GLOBAL	PERFM	ACCTG	STAT	AUDIT	MONITOR	ALL
SMF	√	√	√	√	√	√	√
GTF	√	√	√	√	√	√	√
SRV (not valid for DB2 V2.3 or higher)	√	√	√	√	√	√	√
RES	√						√
OPn	√	√	√	√	√	√	√
Plans	√	√	√		√	√	√
Authids	√	√	√		√	√	√
Classes	√	√	√	√	√	√	√
RMIDs	√	√				√	√
Trace IDs	√	√	√	√	√	√	√
Location	√	√	√		√	√	√

Introduction to DB2 Commands for Utilities

Overview

This unit covers options that manage utilities, required authority for each option, and the values required to use the option. (With these two options, the required values are the same.)

Options that manage utilities

Use these !DB/WORKBENCH options to manage utilities:

- Display Utility Status (DU)
 - Stop a Utility (PU)
-

Authority required for Display Utility Status

No authority is required.

Authority required for Stop a Utility

If you want to use the option for Stop a Utility, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
 - Implied authority, because you started the utility
-

Values required for Display Utility Status and Stop a Utility

To access the DB2 commands available with the options, you provide a value for this field.

Field Name	Description
Utility ID	"*", full ID, or id*

Introduction to DB2 Commands for the Archive Log

Overview

This unit contains introductory information on the !DB/WORKBENCH options for the DB2 commands for the Archive Log. It also provides information you will need to use the options.

Options that manage the Archive Log

Use these options to manage the archive log:

- Archive Log (AR)
 - Display Archive (DA)
 - Set Archive (SA)
-

Authority required for Archive Log

If you want to use the option for Archive Log, you must have one of these authorizations:

- SYSADM, SYSCTRL, or Installation SYSOPR authority
 - ARCHIVE privilege
-

Values required for Archive Log

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Quiesce	"Y" to Quiesce
TIME	Represents the quiesce time in seconds (Required only if Quiesce = Y)
WAIT	"N"/"Y" (Required only if Quiesce = Y)

Authority required for Display Archive

If you want to use the option for Display Archive, you must have one of these authorizations:

- SYSADM, SYSCTRL, or Installation SYSOPR authority
- ARCHIVE or DISPLAY privilege

Values required for Display Archive

You do not need to provide values to begin to use the Display Archive option.

Authority required for Set Archive

If you want to use the option for Set Archive, you must have one of these authorizations:

- SYSADM, SYSCTRL, or Installation SYSOPR authority
- ARCHIVE privilege

Values required for Set Archive

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Default	Y - Reset Count and Time values specified at DB2 Installation time
Count	1–99 - Maximum number of Tape Units to be dedicated for reading Archive Logs
Time	Minimum Inactivity time before Archive Read Tape is deallocated Minutes (0–1439) Seconds (1–59)
Nolimit	Y - Permanent Allocation of all tape units specified by Count

Introduction to DB2 Commands for Buffer Pools

Overview

This unit covers options that manage buffer pools, required authority for each option, and the values required to use the option.

Options that manage buffer pools

Use these !DB/WORKBENCH options to manage buffer pools:

- Alter Bufferpool (AB)
 - Alter Group Bufferpool (AGBP)
 - Display Bufferpool (DB)
 - Display Group Bufferpool (DGBP)
-

Authority required for Alter Bufferpool

If you want to use the option for Alter Bufferpool, you must have one of these authorizations:

- SYSADM authority
- SYSCTRL authority
- SYSOPR authority

Values required for Alter Bufferpool

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Bufferpool Name	The name of the buffer pool
Virtual Buffer Pool Size	An integer that specifies the number of buffers for the virtual buffer pool size
Hiperpool Size (HPSIZE)	An integer that specifies the number of buffers for the hiperpool
Hiperpool Sequential Steal Threshold (HPSEQT)	A value between 0–100. The value is a percentage that specifies the sequential steal threshold for the hiperpool
Virtual Buffer Pool Sequential Steal Threshold (VPSEQT)	A value between 0–100. The value is a percentage that specifies the sequential steal threshold for the virtual buffer pool.
Virtual Buffer Pool Parallel Sequential Threshold (VPPSEQT)	A value between 0–100. The value is a percentage that specifies the parallel sequential threshold for the virtual buffer pool.
Deferred Write Threshold (VDWQT)	A value between 0–90. The value specifies the vertical deferred write threshold for the virtual buffer pool.
CASTOUT (Y/N)	Y or N for the CASTOUT attribute

Authority required for Alter Group Bufferpool

If you want to use the option for Alter Group Bufferpool, you must have one of these authorizations:

- SYSADM authority
- SYSCTRL authority
- SYSOPR authority

Values required for Alter Bufferpool

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Group Bufferpool Name	The name of the of the group buffer pool
Data Sharing Group Name	The DB2 data sharing group name used to qualify the coupling facility structure name
Entry/Page Ratio	The desired ratio of directory entries to data pages
Class Castout Threshold	The percentage of data entries in use in the group buffer pool before class castout is started
Group BP Castout Threshold	The percentage of data pages in use in the group buffer pool before data pages are castout to DASD
Group BP Checkpoint Interval	The number of minutes between group buffer pool checkpoints
Automatic Recovery (Y/N)	Y or N for DB2 automatic recovery

Authority required for Display Bufferpool

If you want to use the option for Display Bufferpool, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- DISPLAY privilege

Values required for Display Bufferpool

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Bufferpool Name	Active, "*", or BP name
Detail	Level of Reporting I Interval - Since last incremental display * Accumulated - Since last BP activation Default Summary Reporting
List	List Open Index and Table spaces related to BP A Active - List spaces Currently in use * Accumulated - Since BP activation Default Summary Reporting
LSTATS	Incremental List Open spaces Y Yes - Displayed stats are incremental from last DISPLAY Default None

Authority required for Display Group Bufferpool

If you want to use the option for Display Group Bufferpool, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- DISPLAY privilege

Values required for Display Group Bufferpool

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Group Bufferpool Name	The name of the group buffer pool
Data Sharing Group Name	The of the DB2 data sharing group
Connection Type	<p>Displays the group buffer pool information according to connection type</p> <p>G Group buffer pools that are connected to any member of the data sharing group</p> <p>M Group buffer pools that are connected to the member of the data sharing group issuing the command.</p> <p>* All group buffer pools</p>
Connection List Report (Y/N)	Y or N for a report of connection names and status for subsystems currently connected to the group buffer pools

Values required for Display Group Bufferpool (Continued)

Field Name	Description
Member Detail Report	The type of member detail report * Accumulated statistics since the member first connected to the group buffer pool I Incremental statistics since the last Incremental member detail report
Group Detail Report	The type of group detail report * Accumulated statistics since the group buffer pool was allocated or re-allocated I Incremental statistics since the last Incremental group detail report

Introduction to DB2 Commands to Recover BSDS

Overview

This unit contains introductory information on the !DB/WORKBENCH options for the DB2 commands for bootstrap dataset (BSDS). It provides information you will need to use the options.

Background about RECOVER BSDS

When you recover a bootstrap dataset, you need to perform several preparatory actions using access method services. You delete or rename the invalid BSDS to have the same name as the one that encountered the error. When you use the RECOVER BSDS option, the system copies the valid bootstrap dataset into the new dataset and reestablishes the dual mode of the BSDS.

Authority required for RECOVER BSDS

If you want to use the option for RECOVER BSDS, you must have one of these authorizations:

- SYSADM and SYSCTRL authority
 - BSDS privilege
-

Values required for RECOVER BSDS

No values are required to use the option.

Introduction to DB2 Commands for Data Sharing

Overview

This unit contains introductory information on the !DB/WORKBENCH options for the DB2 commands for data sharing. It provides information you will need to use the options.

Options that manage data sharing

Use these !DB/WORKBENCH options to manage data sharing:

- Display Group (DG)
 - Display Locations (DL)
 - Reset Generic LU (RG)
-

Authority required for Display Group

If you want to use the option for Display Group, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
 - DISPLAY privilege
-

Values required for Display Group

To access the DB2 command available with the option, you provide a value for this field.

Field Name	Description
Detail (Y/N)	Y includes detailed information regarding the parallelism subsystem parameters for active members of the data sharing group Y valid for DB2 Ver. 5.1 or higher

Authority required for Display Locations

If you want to use the option for Display Locations, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- DISPLAY privilege

Values required for Display Locations

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Location Name	* - Default
Detail (Y/N)	N - Default

Authority required for Reset Generic LU

If you want to use the option for Reset Generic LU, you must have one of these authorizations:

- SYSADM, SYSCTRL, or SYSOPR authority
- DISPLAY privilege

Values required for Display Locations

To access the DB2 command available with the option, you provide values for these fields. The fields are displayed in the order they appear on the panel.

Field Name	Description
Network LU Name	The real VTAM LU name of the partner whose dat will be purged * will cause all partners information to be purged
Network ID	The real VTAMname of the network that the LU name resides in

Managing the System

Chapter 3. Using Authorizations to Manage the System

Introduction

This chapter provides an introduction to the display of authorizations available with the !DB/WORKBENCH Authorizations Menu and includes information on display of authorizations from an object list panel. Each chapter in this section is independent and can be used without referencing the others.

Chapter contents

Resources for Learning about Topics in This Section	98
Overview of the Chapter	99
Resources for Learning about Authorizations	101
Introduction to Methods to Display Authorizations	102
Accessing the Authorizations Menu	104
Displaying Authorizations from the Menu	105
Displaying Authorizations from an Object List Panel	106
Controlling the Display of Authorizations	107
Tasks Commonly Performed with Authorizations	108

Resources for Learning about Topics in This Section

Overview

This unit lists resources available for learning how !DB/WORKBENCH can be used to manage your DB2 system. It includes information about the resources in this guide and in !DB/WORKBENCH online Help.

Resources in this guide

This section contains an introduction, instructions, and prerequisites for using the facilities in !DB/WORKBENCH for system administration. In this section, each chapter is independent; you do not need to read one chapter before using the other.

Information You Want	Resource Available
Overview of management of authorizations	This chapter
Overview of the DASD estimators for table spaces and index spaces	The chapter “Using the DASD Estimators” on page 111
Overview of how to display ZPARM parameters and to use functions from the DSNZPARM Menu	The chapter “Using the DSNZPARMs Facility” on page 123
Overview of the functions available from the Housekeeping Menu, including using the Restore (Lazarus) facility	The chapter “!DB/WORKBENCH Housekeeping Options” on page 139

Overview of the Chapter

Background about authorizations and !DB/WORKBENCH

With any database management system, maintaining the security of the system is vital to the successful use of the system. Within DB2, authorization can be given to different users. The entire authority falls under the control of and cascades down from the systems administrator (SYSADM).

DB2 authorization requires that all users be authorized to do whatever they need to do. Every DB2 capability performed by a user, represented by an authorization ID, must be authorized, either explicitly or implicitly. For example, activities such as creating a table space or database, creating views, dropping a table, and executing DB2 operator commands need authorization.

As you implement authorization, you develop a complex set of information within the system. You will need to know about this information to successfully implement, monitor, and control your DB2 system. The system stores this information in authorization tables in the DB2 Catalog, but DB2 does not provide automated facilities to look at what you have specified.

With !DB/WORKBENCH, you can access the information using the !DB/WORKBENCH panels instead of writing SQL statements or using SPUFI to get the information you need.

Background about enforcement by !DB/WORKBENCH

!DB/WORKBENCH does not enforce any limitation on a user action that DB2 does not enforce. In this way, a user without full authority to perform a task can begin preliminary steps and then let another user with authority execute the task.

Organization of this chapter

This chapter provides an introduction to the display of the authorizations available with !DB/WORKBENCH. It includes display of authorizations by using the Authorizations Menu or by using the the Auths (Z) select from many object list panels.

Organization of information and your needs

Review the chart to select information appropriate to the task that you want to perform.

	Resources	Introduction	Accessing	Menu	Object list panels	Controlling	Using
	p.101	p.102	p.104	p.105	p.106	p.107	p.108
Review resources for learning about details on authorizations using online Help	√						
Compare the two methods available for displaying authorizations		√					
Display and use authorizations that you access from the Authorizations Menu	√	√	√	√		√	√
Display and use authorizations that you access from an object list panel	√	√			√	√	√

Resources for Learning about Authorizations

Overview

This unit covers the resources available for authorizations through the online Help facility.

Online Help and authorizations

The online Help available with authorizations includes detailed explanations of the fields on the panels, as well as instructions for using commands and selects available from a specific panel.

```

Cmd ==>
----- HELP FOR TABLES AND VIEWS -----
Press END to return, ENTER to continue
or TAB to highlighted text and press ENTER to display more information.

Description      Use this panel to review the authorizations for tables and
                  views. Each table or view name is prefixed by the creator ID.
                  You can only review authorization for objects you are
                  authorized to view.

Filtering        You can use the authorization FILTERS to limit the objects
                  displayed (and included in output grants).

Detail
Information      You can review the DETAIL AUTHORIZATION information for an
                  object by typing S in front of the object.

Column
Information      You can display all the COLUMN AUTHORIZATIONS for a table or
                  view by typing C in front of the object.

-----
GLOSSARY  INDEX  COMMANDS  FILTERS  WINDOWS  SELECTS  MORE->

```

Introduction to Methods to Display Authorizations

Overview

This unit lists the methods to display authorizations, the authorizations panel, and the differences between the types of authorizations.

Methods to display authorizations

!DB/WORKBENCH provides two methods for displaying authorizations. The chart shows the differences between the methods.

Information You Want to Display	Method to Use	Authority Required
All authorizations for the object	Authorizations Menu	Yes
All authorizations that you are authorized to see	Auths (Z) select on an object list panel	No

The panels that display authorizations

If you are displaying the authorizations for the same type of object from the Authorizations Menu or from the object list panels, !DB/WORKBENCH displays the same panel in either case. The difference is in the quantity of authorizations that !DB/WORKBENCH displays with the two methods.

For example, if you are the database administrator and select the Databases option from the Authorizations Menu, !DB/WORKBENCH displays all authorizations for databases. If you are authorized to use a single database only and you use the Auths (Z) select on the Databases panel, !DB/WORKBENCH displays the same panel but lists only authorizations for the single database.

Differences in types of authorizations available

The chart indicates what authorizations are available from the Authorizations Menu and by using the Auths (Z) select on an object list panel.

Object	Available from the Menu	Available from the Panel
Buffer pools	√	
Collections	√	
Columns	√	
Databases	√	√
Grantees (for the number of grants to each grantee by type or grant)	√	
Packages	√	√
Plans	√	√
Storage Groups	√	√
Table Spaces	√	√
Tables	√	√
Users (for a system-level authorization, such as being a SYSADM)	√	√
Views	√	

Accessing the Authorizations Menu

Overview

This unit contains instructions for accessing the Authorizations Menu.

Accessing the Authorizations Menu

If authorized, you can access the Authorizations Menu by following these steps.

Step	Action
1	On the !DB/WORKBENCH Primary Menu, type Z in the Option field.
2	<p>Press Enter.</p> <p>Result: The system displays the Authorizations Menu.</p> <div data-bbox="573 884 1248 1297" style="border: 1px solid black; padding: 10px;"> <pre> AUTHORIZATIONS ----- DB/WORKBENCH -- DB2=DB2 ----- Cmd ==>> 99/12/08 13:24 D B 2 O B J E C T A U T H O R I Z A T I O N Option ==>> 1 D DATABASES Display Auths on Databases. 2 S TABLESPACES Display Auths over Tablespaces. 3 T TABLES Display Auths over Tables. 4 V VIEWS Display Auths over Views. 5 C COLUMNS Display Auths over Columns. 6 P PLANS Display Auths over Plans. 7 B POOLS Display Auths over Bufferpools. 8 G STOGROUPS Display Auths over Storage Groups. 9 U USERS Display User Authorizations. 10 L COLLECTIONS Display Auths over Collections. 11 K PACKAGES Display Auths over Packages. 12 E GRANTEE Display List of Grantees. </pre> </div>

Displaying Authorizations from the Menu

Overview

This unit shows how you can display authorizations by using the menu or using the quick access method.

Displaying authorizations using the Authorizations menu

If you are authorized to use the Authorizations menu, follow these steps to display authorizations.

Step	Action
1	On the Authorizations Menu, type the letter or number for the option that you want in the Select Option field. Example: If you want to display authorizations for databases (which you can select with a 1 or D), you type 1 or D .
2	Press Enter. Example: The system displays all authorizations for the object you selected. In this case, the system displays all authorizations for databases.

Displaying authorizations from the !DB/WORKBENCH Primary Menu (quick access)

If you are authorized to use the Authorizations Menu, follow these steps to go directly to the object list panel showing authorizations.

Step	Action
1	On the !DB/WORKBENCH Primary Menu, type Z followed by a period (.), followed by the character of the selection (on the Authorizations Menu) that you want in the Select Option field. Example: If you want to display authorizations for databases (which you can select with 1 or D), you type Z.1 or Z.D .
2	Press Enter. Example: The system displays all authorizations for the object you selected. In this case, the system displays all authorizations for databases.

Displaying Authorizations from an Object List Panel

Overview

This unit contains instructions for displaying authorizations from an object list panel. No authorization is required for this procedure.

Displaying authorizations from an object list panel

If the Auths (Z) select is available on the object list panel, follow these steps to display the authorizations for the object that you are allowed to see.

Step	Action
1	On the object list panel, type Z in the Select field. Example: On the Database panel, type Z in the Select field.
2	Press Enter. Result: The system displays only the authorizations that you are authorized to see for the object you selected.

Controlling the Display of Authorizations

Overview

This unit lists the resources for learning how to use filters to control the display of authorizations and shows the differences between filters for authorizations and filters for other objects.

Using filters to control the display

Although filters for authorizations operate like filters for other objects, differences exist. For example, filters for authorizations are controlled separately from filters for other objects. The chart lists resources available on filtering with authorizations.

Information You Want	Source to Use
How filtering with authorizations differs from other filters	The chapter “Filtering Object Lists” in <i>!DB/WORKBENCH Basic User's Guide</i>
Methods available for filtering authorizations	The chapter “Methods for Filters—Authorizations” in <i>!DB/WORKBENCH Reference Manual</i>

If using filters for authorizations for tables

When used on the panel displaying authorizations for tables, the DB (database) filter does not filter on the same database name used with the other panels. On tables, the DB filter uses the database name that is in the DB2 Catalog as the Grantor DBADM database and not the database that owns the table.

Tasks Commonly Performed with Authorizations

Overview

This unit discusses how the panels that display authorizations are organized, what commands and selects are available on the panels, and examples of commands and selects that can be used with tasks for authorizations.

Background about the organization of the panels displaying authorizations

Panels displaying authorizations for objects (such as databases, tables, users, and views) include a series of columns with abbreviated names. The columns provide details about specific authorizations. !DB/WORKBENCH provides a select so you can display the full names for these authorizations.

For example, on the Databases panel, the series of columns with details includes the CT column. The CT column can include these values: G (for Grant), Y (for Yes), or blank (for No). If you use the S select, !DB/WORKBENCH displays a panel that lists the full names for the columns and for the values in the columns. For example, the system displays detailed information about a CT column with a value of G in this way:

```
CREATE TABLES ----- GRANT
```

Availability of commands and selects on panels

The only commands and selects that are available on the panel are those listed on the panel.

Commands and selects available for tasks with authorizations

The chart shows examples of commands and selects that are commonly available with authorizations. The actual characters used with the selects and the commands, however, may vary from panel to panel. Review your panel for the options available and the letters to use. (For information about GRANTS, see the chapter “Methods for Generating Statements for Objects” in *!DB/WORKBENCH Basic User's Guide*.)

Task You Want to Perform	Method	Example
Display explanation of the full names for the details about authorizations for a specific object or provide other details, such as the date-time order of grants	Select	S
Generate GRANT statements for a specific object	Select	O
Generate GRANT statements for all objects displayed on the panel	Command	OUT

Tasks Commonly Performed with Authorizations

Introduction

This chapter contains prerequisites and procedures for using the DASD estimators for table spaces and index spaces.

Chapter contents

Overview of the Chapter	112
Resources for Using the Estimators	113
Accessing the Estimators	114
Functions Available with the Estimators	116
Prerequisites for Estimating DASD Requirements	118
Values Reported by the Estimator	120

Overview of the Chapter

Background about the DASD estimators and !DB/WORKBENCH

!DB/WORKBENCH provides DASD space estimators for you to use with table spaces and index spaces. You can enter values when creating a new table space or index space, and the system calculates the effects of the changes and displays the results.

Organization of the chapter and your needs

Review the chart to select information appropriate to the task that you want to perform.

	Resources	Accessing	Functions	Prerequisites	Report
	p.113	p.114	p.116	p.118	p.120
Review the resources that are available for using the DASD estimators	√				
Using the estimator for table spaces		√	√	√	√
Using the estimator for index spaces		√	√	√	√

Resources for Using the Estimators

Overview

This unit covers the resources available that will help you use the estimators for table spaces and index spaces.

Panels and the estimators

The !DB/WORKBENCH panels for the Table Space Estimator and the panel for the Index Space Estimator include directions for entry. For example, the panel for the Index Space Estimator lists the possible values (1 to 255) and the definition *Length of Key field* for the KEYLENG field.

Online Help and the estimators

If you want additional instructions about the values for the estimators, you can use the online Help that is available for the estimator panels. For example, if you want additional information about the values on the Index Space Estimator panel, type **HELP** on the command line and press Enter, or press the appropriate function key. The system displays this Help panel.

```

Cmd ==>
----- HELP FOR INDEX SPACE ESTIMATOR -----
Press END to return, ENTER to continue
or TAB to highlighted text and press ENTER to display more information.

Description      Use this panel to calculate the size for a new index space.

KEYLENG          Specify the length in characters of the key field in this
                  index. Type a value in the range 1-255.

ROWS             Type the estimated number of table rows to be indexed.

PCTFREE          Specify the percentage of free space in each page. Type a
                  value in the range 0-99.

FREEPAGE         Specify the number of pages between free pages. Type a value
                  in the range 0-255.

-----
GLOSSARY  INDEX  COMMANDS  FILTERS  WINDOWS  SELECTS  MORE->

```

Accessing the Estimators

Overview

You use the same procedure to access the estimators for table spaces and index spaces. This unit shows the procedures for accessing the estimators.

Procedure for accessing the estimators from the !DB/WORKBENCH Primary Menu

Use these steps to access the table space or the index space estimator from the !DB/WORKBENCH Primary Menu.

Step	Action
1	<p>On the !DB/WORKBENCH Primary Menu, type the appropriate value in the Options field.</p> <ul style="list-style-type: none"> ● For estimated values for table spaces, type ES ● For estimated values for index spaces, type EX
2	<p>Press Enter.</p> <p>Result: The system displays the panel for you to use to estimate the DASD requirements for the table space or the index space.</p> <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> TABLESPACE ESTIMATOR ----- DB/WORKBENCH --- DB2=DB31----- Cmd ==> 99/07/14 09:49 PAGESIZE ==> Y Y 4K Page N 32K Page COMPRESSION ==> 0 (0-99) Percent compression PCTFREE ==> 10 (0-99) Percent freespace per page FREEPAGE ==> 0 (0-255) Pages between free pages FUDGE ==> 10 (0-99) Percentage to add to estimates SEGSIZE ==> 0 (0-64) Size in increments of 4 UNIT TYPE ==> T T Tracks C Cylinders AVG ROW SIZE NUMBER OF ROWS PER PAGE PAGES OCCUPIED 3380 3390 ----- ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 ==> 0 ==> 0 0 0 0 0 TOTAL SIZE - - - 0 0 0 Fill in values, then press ENTER. Press END when done. </pre> </div>

Accessing and using the estimators when creating objects

You can access the estimators when you are creating a new table space or a new index space using the options available from the Create Objects Menu. This menu is accessed when you type **CR** on the Primary Menu. Use these steps to access and use the table space estimator or the index space estimator from the Create Tablespace panel or the Create Index panel, respectively.

Step	Action
1	On the !DB/WORKBENCH panel listing the values for the new table space or index space, type ? in the Priqty field. This field will display on a panel several panels into the create process.
2	Press Enter. Result: The system displays the panel for you to use to estimate DASD requirements for the table space or index space.
3	Type new values for any fields you want to modify.
4	Press Enter. Result: The system displays the results of the estimates on the lower half of the panel.
5	Type END on the command line and press Enter, or press the appropriate function key. Result: The system displays the results of your calculations in the Priqty field on the panel that you were using to create the table space or index space.

Functions Available with the Estimators

Overview

This unit describes how the estimators operate and the types of DASD estimators available in !DB/WORKBENCH.

How the estimators operate

To use the estimators, you provide values in the fields listed on the upper half of the panel. (With table spaces, you also provide essential values for each table in the table space you are estimating.)

When you press Enter, the estimator calculates new values and displays the estimates in the lower half of the panel. You can continue to modify values and press Enter to display new estimates until you exit the estimator. Use these estimates to plan for storage needs when creating new objects.

Types of estimators available with !DB/WORKBENCH

With !DB/WORKBENCH, there are two estimators available to estimate DASD requirements:

- Table Space Estimator
 - Index Space Estimator
-

Using the estimators

Follow these steps to estimate values for a table space or index space. (You can repeat this procedure to prepare estimates using different values.)

Step	Action
1	Type new values for any field you want to modify.
2	Press Enter. Result: The system displays the results of the estimates or the lower half of the panel.

Exiting the estimators

When you have completed the estimates for the table space or index space, you can exit the estimator by typing **END** on the command line and pressing Enter, or by pressing the appropriate function key.

Prerequisites for Estimating DASD Requirements

Overview

This unit gives background information about the fudge factor and the unit value. It also lists the prerequisites for estimating DASD requirements for table spaces and index spaces.

Background for the fudge factor

The DASD estimators for table spaces and index spaces allow you to use a fudge factor. The fudge factor is the estimated percentage of additional free space. You can use any value from 0 to 99; the default is 10.

The value you use as a fudge factor depends on the method you plan to use to build the table or a new index on a table. The chart provides guidelines on how the method you use determines the appropriate value to use for either a new table space containing the table or a new index on a table.

Method to Build the Table	Value to Use
LOAD	A low value
INSERT	A high value (such as 50)

Background for the Unit value

The DASD estimators for table spaces and index spaces can display the output in the unit you prefer: in tracks (T), or in cylinders (C).

Prerequisites common to estimating table spaces and index spaces

These prerequisites are common to estimates for both table spaces and index spaces:

- The percent of free space per page (PCTFREE)
- The pages between free pages (FREEpage)
- The fudge factor
- The Unit value

Additional prerequisites for table spaces

These prerequisites are specific to estimates for table spaces:

- The pages between free pages (PGSIZE)
- Whether data compression is in effect (COMPRESS) - used only with DB2 V3.1 or greater
- The SEGsize (if a segmented table space)
- Average size of a row in each table in the table space (which can be calculated using the DSN1COMP utility)
- Estimated number of rows in each table in the table space

Background about key length and indexes

If you are using the Estimator when creating a table and index, you do not need to provide the value for the KEYLENG field on the panel displayed when you use the Estimator. The system uses the value from the creation of the table as the default value for the estimates.

Additional prerequisites for index spaces

These prerequisites are specific to estimates for index spaces.

- The KEYLENG value (if you have not been using the CR option from the !DB/WORKBENCH Primary Menu)
- The number of rows in the table to be indexed
- Whether the index is unique or allows duplicate keys (If the index allows duplicate keys, you also must know the number of keys using the DISTINCT or RECS/KEY fields.)

For more information about values for prerequisites

The most immediate source for the ranges of values for the fields is the online Help for the estimators. You also can use the values in the chapter “Values Used When Creating or Altering Objects” in *!DB/WORKBENCH Reference Manual*.

Values Reported by the Estimator

Overview

This unit lists the values calculated and reported by the table space estimator and the index space estimator. It also gives information about using the online Help to get more detailed information about the values.

Values reported by the table space estimator

The table space estimator reports these values:

- The average number of rows per page
 - The number of pages required for the estimated number of rows
 - The number of 3380 tracks or cylinders to contain the table in the table space
 - The number of 3390 tracks or cylinders to contain the table in the table space
 - Total size (the value used in the Priqty field when creating a table space)
-

Values reported by the index space estimator

The index space estimator reports these values:

- The number of subpages
- The subpage size
- The number of rows that fit in each subpage
- The number of levels in the index
- The total number of pages required for the index
- The number of 3380 tracks or cylinders to hold the index
- The number of 3390 tracks or cylinders to hold the index

Online Help and the field definitions

If you want detailed field definitions for the values for the estimators, you can use the online Help that is available for the estimator panels. For example, if you want additional information about the values on the Table Space Estimator panel, type **HELP** on the command line and press Enter, or press the appropriate function key. The system displays this Help panel.

```

Cmd ==>
----- HELP FOR TABLE SPACE ESTIMATOR -----
Press END to return, ENTER to continue
or TAB to highlighted text and press ENTER to display more information.

AVG ROW SIZE      NUMBER OF ROWS      AVG ROWS      PAGES      3380      3390
                   PER PAGE      OCCUPIED      TRACKS      TRACKS
-----
=== |Number of rows that fit |-----> 7      110      11      10
=== |on a page after DB2    |-----> 8      8      >1      1
=== |overhead.             |-----> 165 ] ----->14
===>|----->0             |-----> 0 ] ] ----->0
===>0                 ===> 0             0 ] ] ----->0
                   ===> 0             0 ] ] ----->0
                   ===> 0             0 ] ] ----->0
                   ===> 0             0 ] ] ----->0
|Primary space | ===> 0             0 |Number of tracks or |
|allocation in | ===> 0             0 |cylinders required. |
|K-bytes. Use  | ===> 0             0 |----->0
|this value |--->TOTAL SIZE - | 1,140K -      285      29      24
|in the |----->
|PRIQTY field of the |
|Create Table Spaces panel. |
|----->
-----
GLOSSARY INDEX COMMANDS FILTERS WINDOWS SELECTS <-PREV

```

Values Reported by the Estimator

Chapter 5. Using the DSNZPARMs Facility

Introduction

!DB/WORKBENCH provides a facility to view and modify the parameters in the system initialization module. This chapter discusses the functions available with the !DB/WORKBENCH DSNZPARMs facility and gives procedures to access the DSNZPARMs menu and to select the options.

Chapter contents

Overview of the Chapter	124
Functions Available with the DSNZPARMs Facility	126
Accessing the DSNZPARMs Menu	128
Selecting Options from the DSNZPARMs Menu	130
Displaying and Editing the DSNZPARMs Parameters	131
Generating Macros from the DSNZPARMs Menu	134
Migrating DB2 Parameters	136
Printing DB2 Parameters	137

Overview of the Chapter

Background about DSNZPARMs and !DB/WORKBENCH

DSNZPARM is the name for a module that contains the System Initialization parameters. The module contains the parameters used to initialize DB2. The !DB/WORKBENCH facility provides an easy way to view the DSNZPARMs. The Online Help in !DB/WORKBENCH contains the definitions of the ZPARAM keywords and will save you the time used in checking reference manuals for the meaning of the keyword parameters.

Terminology in this chapter

The name of the actual module used to initialize DB2 can be any name chosen by your installation. Throughout this chapter, the module itself is referred to as “DSNZPARM,” and the parameters themselves are called “ZPARAMs.” The ZPARAMs panels identify the current values of the parameters and give their definitions and value options.

Organization of this chapter

This chapter focuses on the functions available with the DSNZPARMs facility. It provides instructions for accessing the menu and selecting options.

Organization of information and your needs

Review the chart to select information appropriate to the task that you want to perform.

	Functions	Accessing	Selecting	Displaying	Generating	Migrating	Printing
	p.126	p.128	p.130	p.131	p.134	p.136	p.137
Review functions available in the DSNZPARMs facility	√						
Display DSNZPARMs	√	√	√	√			
Generate DSNZPARMs macros	√	√	√		√		
Migrate DSNZPARMs parameters	√	√	√			√	
Print DSNZPARMs parameters	√	√	√				√

Functions Available with the DSNZPARMs Facility

Overview

This unit shows how you can use the DSNZPARM facility to display system, installation, and restart information. It also gives a list of tasks that can be performed with the facility and the option to select for each task.

Displaying system information with the DSNZPARMs facility

With !DB/WORKBENCH, you can display and change the DB2 macros that the system uses to build the DSNZPARM module. (The name of these macros can vary on your system.)

Information You Want to Display	Relationship to DB2 Macro	Option to Select
Archive Dataset Parameters	DSN6ARVP	ARVP
DB2 Logging Parameters	DSN6LOGP	LOGP
System Parameters	DSN6SYSP	SYSP
System Authorization Parameters	DSN6SPRM (and DSN6ENV)	SPRM
Data Sharing Facility Parameters	DSN6GRP	GRP
Distributed Data Facility Parameters	DSN6FAC	FAC

Displaying installation and restart information

Using the DSNZPARMs facility, you can display and change information about the DB2 restart parameters for your system and the installation parameters for your DB2 system.

Information You Want to Display	Option to Select
DB2 Restart Parameters (a list of any databases that are defined for automatic start at DB2 initialization)	RESTART
Installation Parameters (the release and maintenance levels of your DB2 system)	DSNHDECP

Performing system tasks with the DSNZPARMs facility

With the DSNZPARMs facility, you can generate DB2 macros, migrate your current setting of the ZPARMs to a new version of DB2, and generate a report that shows all the ZPARMs for your DB2 system.

Task You Want to Perform	Option to Select
Generate the DB2 macros used to assemble the DSNZPARMs and write them to your !DB/WORKBENCH user PDS	MACROS
Migrate your current settings of the ZPARMs values to a new version of DB2 (The option uses the library name and member you specify and stores your current system ZPARM values in a format that can be used by the IBM install panels and CLIST.)	MIGRATE
Write to a JES printer a report showing the current ZPARM values.	PRINT REPT

Accessing the DSNZPARMs Menu

Overview

This unit shows the procedures used to access the DSNZPARMs menu.

Authorization to display ZPARMs values

Your installation may have elected to restrict the display of the values for the ZPARMs. If so, you may not be able to access these panels. When you use ZP from the !DB/WORKBENCH Primary Menu, you receive an error message indicating that this option is not available to you.

Procedure for accessing the DSNZPARMs menu

Use these steps to access the DSNZPARMs menu from the !DB/WORKBENCH Primary Menu.

Step	Action
1	On the !DB/WORKBENCH Primary menu, type ZP in the Option field.
2	<p>Press Enter.</p> <p>Result: The system displays the menu of options available for DSNZPARMs.</p> <div data-bbox="573 1173 1248 1589" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> ZPARMS OPTIONS ----- DB/WORKBENCH -- DB2=D31A----- Cmd ==> 99/12/19 07:41 Option ==> 1 SPRM System Authorization Parameters. 2 ARVP Archive Data Set Parameters. 3 LOGP DB2 Logging Parameters. 4 SYSP System Parameters. 5 GRP Data Sharing. 6 FAC DDF Facility. 7 DSNHDECP Installation Parameters. 8 RESTART DB2 Restart Parameters. 9 MACROS Generate ZPARMS Macros. 10 MIGRATE Update DB2 Migration Parameters. 11 PRINT REPT Print ZPARMS report on JES printer. </pre> </div>

Procedure for accessing an option on the DSNZPARMs menu (quick access)

Use these steps to access an option on the DSNZPARMs menu from the !DB/WORKBENCH Primary Menu.

Step	Action
1	On the !DB/WORKBENCH Primary Menu, type ZP followed by a period (.), and followed by the number of the option you want. Example: If you want to access the System Authorization parameter, type ZP.1 in the Option field.
2	Press Enter. Example: The system displays the System Authorization parameters.

Selecting Options from the DSNZPARMs Menu

Overview

This unit shows you how to select options from the DSNZPARMs menu.

Procedure to select options

Use these steps to select the options from the DSNZPARMs menu.

Step	Action
1	On the DSNZPARMs menu, type the number for the option that you want in the Select Option field.
2	Press Enter. Result: The system displays the data you requested or performs the task.

Displaying and Editing the DSNZPARMs Parameters

Overview

This unit lists the types of ZPARM data that are available for display and the options you have to display or change the data.

DSNZPARM parameter types

The !DB/WORKBENCH DSNZPARMs facility contains the following groups of ZPARM parameters that you can view and edit:

- System Authorization Parameters
 - Archive Dataset Parameters
 - DB2 Logging Parameters
 - System Parameters
 - Data Sharing Parameters
 - DDF Facility Parameters
 - Installation Parameters
 - DB2 Restart Parameters
-

Displaying the !DB/WORKBENCH DSNZPARM panels

To display the !DB/WORKBENCH DSNZPARM panels for a particular group of ZPARM parameters, type the option number of the parameter group on the ZPARMS Options panel, and press Enter.

DSNZPARMs panels

The DSNZPARM panels identify the library that is the source of the ZPARM display. The display of data available from DSNZPARMs panels includes:

- The name of the ZPARM parameter
- The current value for the ZPARM parameter
- A brief description of the ZPARM parameter

For example, the !DB/WORKBENCH DSNZPARMs panel shown below displays the field *MVS* with a value *XA* and the description *System environment—always “XA”*.

```

DISPLAY ZPARMS DSN6SPRM--- DB/WORKBENCH --- DB2=DB2 ----- LINE 1 OF 145
Cmd ==>                               Scroll ==> PAGE
          Field values from PP.DB2.V2R3.DSNLOAD(DSNZPARM)

Enter ? over data to display help. Zap-only fields begin with #. If a field
is in error, clear the field to reset it.           Scroll through display.
-----
Field          Value          Description
-----
MVS            ==> XA           System environment - always "XA".
CATALOG        ==> DSNC 0     VSAM ICF Catalog name that owns the system dsets.
SYSADM         ==> SYSADM     System administrator ID.
SYSADM2        ==> SYSADM     System administrator ID for secondary support.
SYSOPR1        ==> SYSOPR     Console operator authorization ID.
SYSOPR2        ==> SYSOPR     Console operator authorization ID.

AUTH           ==> NO         "Y" or "N" - Activate DB2 authorization?
DECDIV3        ==> NO         "Y" or "N" - Zero divide option-don't round cents?
DBCHK          ==> NO         "Y" or "N" - Expanded DB checking? (obsolete)
DEFLTID        ==> IBMUSER    System default userid ID.
RECALL         ==> YES        "Y" or "N" - HSM Auto-Recall? (V2.1)
RECALLD        ==> 120        Seconds to wait for Auto-recall.(V2.1)
#RECALLT       ==> 8          Number of recall service tasks.
    
```

Options available on the panels for displaying or changing data

These options are available for displaying or changing data on the panels:

- Scroll through the DSNZPARMs panels to see additional data.
- Type a question mark (?) in the field showing the current value for the ZPARM parameter, and press Enter to display detailed information on the Help panels.
- Type a new value in the field, and press Enter to save the changes. (You can later generate the macros to implement the changes by selecting the option for Generate ZPARMs Macros.)

Note: ZPARM parameters that are preceded by a pound sign (#) cannot be changed using the DSNZPARMs facility. These Parameters can only be changed with a ZAP.

Generating Macros from the DSNZPARMs Menu

Overview

You can generate or edit DB2 macros using the DSNZPARM facility. These macros are used to assemble the DSNZPARMs and write them to your !DB/WORKBENCH user PDS. This unit gives you the prerequisites and procedure for generating and editing macros.

Prerequisites for using the Macro option

Before using the macro option, you need to perform these tasks:

- Change any values you need to change on the ZPARMs panel.
 - Confirm that the model JCL members exist:
 - JCLZPRM1 that contains the JOB card and instream PROC
 - JCLZPRM2 that contains the linkedit statements
 - If desired, add these members to your installation requirements.
-

Changes that require the generating of new macros

You must generate new macros to implement the following DSNZPARMs changes:

- Changes to a value on one or more DSNZPARMs panels, such as SPRM, ARVP, LOGP, SYSP, GRP, or FAC
- Changes to the RESTART options

Procedure for accessing the DSNZPARMs macro options

Follow these steps to access the DSNZPARMs macro options from the DSNZPARMs menu.

Step	Action
1	On the DSNZPARMs menu, type 9 in the Option field.
2	<p>Press Enter.</p> <p>Result: The system displays the menu of options available for the DSNZPARMs macros.</p> <div data-bbox="570 642 1248 968" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> PROCESS MEMBER ----- DB/WORKBENCH --- DB2=DB31----- Cmd ==>>> 99/12/04 06:59 The ZPARMS Macro statements have been saved in member USERZ. Option ==>>> 1 SUBMIT Submit the member. 2 EDIT Edit saved statements. 3 DELETE Delete saved statements. 4 PRINT Print the statements member. 5 EXIT Exit now. </pre> </div>

DSNZPARMs macro options

With the DSNZPARMs Macro option, you can submit the member as an MVS job, review and edit the !DB/WORKBENCH user PDS member using ISPF editor, delete the job, or print the named member.

Task You Want to Perform	Option to Select
Generate the DB2 macros used to assemble the DSNZPARMs and write them to your !DB/WORKBENCH user PDS	SUBMIT
Review or edit the the !DB/WORKBENCH user PDS using the ISPF editor	EDIT
Delete the named member from the PDS	DELETE
Print the named member on the JES printer using the options specified in !DB/WORKBENCH Housekeeping.	PRINT

Migrating DB2 Parameters

Overview

You can use the DSNZPARMs facility to migrate your current settings of the ZPARM values to a new version of DB2. This unit gives you the prerequisites and procedure for migrating those parameters.

Prerequisites for using the Macro options

To migrate the DB2 parameters, you need to provide these values:

- The dataset name (written using TSO naming conventions)
- The member name

Procedure for migrating DB2 parameters

Follow these steps to migrate DB2 parameters.

Step	Action
1	On the DSNZPARMs menu, type 10 in the Option field.
2	Press Enter. Result: The system displays a panel to enter the required values.
3	Type the dataset name in the Data Set Name field.
4	Type the name of the member in the Member Name field.
5	Press Enter. Result: The system copies the current ZPARM parameters to the dataset member that you specified.

Printing DB2 Parameters

Overview

You can use the DSNZPARMs facility to write a report to a JES printer. The report lists the current ZPARM values. This unit gives you the prerequisites and the procedure to print this report.

Prerequisites for printing DB2 parameters

To print the DB2 parameters, you need to have previously provided these values:

- The printer class
 - The printer destination
-

Procedure for printing DB2 parameters

Follow these steps to print DB2 parameters.

Step	Action
1	On the DSNZPARMs menu, type 11 in the Option field.
2	Press Enter. Result: The system prints the report showing the current ZPARMs for your DB2 system.

Chapter 6.

!DB/WORKBENCH Housekeeping Options

Introduction

This chapter provides an overview of the options available when you select the Housekeeping option on the !DB/WORKBENCH Primary menu. It identifies prerequisite authority required to change default values, and provides information on changing the default values used in the options that comprise Housekeeping.

Chapter contents

Overview of the Chapter	140
Introduction to !DB/WORKBENCH Housekeeping Options	143
Overview of !DB/WORKBENCH Housekeeping	144
Accessing !DB/WORKBENCH Housekeeping Options	146
Option 1: Profile Data Administration	148
Option 1.0: Global Dataset Information	150
Option 1.1: Global Configuration Information	153
Option 1.2: Global Print Options	156
Option 1.3: DB2 Subsystem-Configuration Information	159
Option 1.4: Catalog Prefixes for the IDF	163
Option 1.5: !DB/WORKBENCH Configuration Information	165
Option 1.6: !DB/WORKBENCH Processing Information	169
Option 2: Update Authorizations Information	172
Option 3: Display Extract History	173
Option 4: Run a New Extract	178
Option 5: Display System PDS Members	180
Option 6: Resurrect Old Catalog Copies	182
Option 7: Display the Activity Log	186
Option 8: Display the !DB/Tools Allocated Profile Datasets	189
Option 9: Set the !DB/Tools Profile Datasets	191

Overview of the Chapter

Background about !DB/WORKBENCH Housekeeping options

Use the !DB/WORKBENCH Housekeeping options to

- verify or update profile data and configuration information, such as the
 - !DB/Tools global dataset information
 - !DB/Tools global print options
 - !DB/Tools global configuration information
 - DB2 configuration information
 - DB2 catalog prefixes for the Interactive Display Facility
 - !DB/WORKBENCH configuration information
 - !DB/WORKBENCH processing information
- update the Authorizations information
- list a history of extracts and display detailed information on a selected extract
- run a new extract
- list the system PDS members (groups) and display or delete information in a selected group
- resurrect a previous (LAZARUS) copy of the catalog
- display or modify the activity log of the log PDS
- display the !DB/Tools profile datasets
- set the !DB/Tools profile datasets

During the installation process, default values are set for Housekeeping options in the !DB/Tools products and in !DB/WORKBENCH. You can change the values to meet your specifications.

Organization of information in this chapter

This chapter focuses on the Housekeeping options available in !DB/WORKBENCH. The information is presented in the same order that the Housekeeping options are presented. It provides general instructions on accessing the Housekeeping options. It includes an overview of the options available and a series of units on the authority required and the values you will need to provide when you use the options.

Additional information about extract process and the STATDBAS table space

This chapter covers the functions that you can use in Housekeeping to manage the extracts that !DB/WORKBENCH uses to provide information quickly. The chapter does not, however, cover the statistical data that the system inserts in the tables in the STATDBAS table space each time you run an extract. This feature allows you to run queries against statistical data in the tables in the STATDBAS table space. (The availability of the data in the tables in the STATDBAS table space depends on steps you have followed during the installation of !DB/WORKBENCH.)

If you want more information about statistical data or about controlling the inserting of statistical data in the tables in the STATDBAS table space, see “Managing Statistics in the STATDBAS Table Space” on page 195.

Additional information related to !DB/WORKBENCH Housekeeping options

Additional information is available for these issues related to !DB/WORKBENCH Housekeeping options:

- For more information about user exits, see “Appendix D. !DB/Tools Access Control Statements” in the *!DB/Tools Installation and Customization Guide*.
- For more information about authorizations, see “Using Authorizations to Manage the System” on page 97.

Organization of the chapter

Review the chart to select information on a specific Housekeeping function. Then locate the page number of the information.

	Introduction to Housekeeping	Overview of Options	Accessing Housekeeping	Global Profile Information	DB2 Subsystem Information	Catalog Prefixes	Configuration Information	Processing Information	Update Authorizations	Extract History	Run New Extract	System PDS Members	Restore Catalog Copies	Activity Log	Display Profile Datasets	Set Profile Datasets
	p.143	p.144	p.146	p.150	p.159	p.163	p.165	p.169	p.172	p.173	p.178	p.180	p.182	p.186	p.189	p.191
Review introductory information about Housekeeping	√															
Review overview of options available using Housekeeping		√														
Access Housekeeping			√													
Use the option for global profile information				√												
Use the option for the specific DB2 subsystem					√											
Use the option for catalog prefixes for IDF						√										
Use the option for !DB/WORKBENCH configuration information							√									
Use the option for !DB/WORKBENCH processing information								√								
Use the option for updating user authorizations									√							
Use the option for extract history										√						
Use the option for running a new DB2 catalog extract											√					
Use the option for displaying system PDS members												√				
Use the option for restoring catalog copies													√			
Use the option for displaying the activity log														√		
Use the option for displaying the !DB/Tools profile datasets															√	
Use the option for setting the !DB/Tools profile datasets																√

Introduction to !DB/WORKBENCH Housekeeping Options

Overview

You perform housekeeping tasks to maintain the integrity of your system. This unit contains information about performing these administrative tasks using !DB/WORKBENCH Housekeeping options. It includes reasons for performing housekeeping tasks.

Reasons for performing housekeeping tasks

The chart lists common reasons for performing housekeeping tasks.

Reasons for Performing Housekeeping Tasks	Example
To compare the difference in the number of objects of the last 5 extracts that were run	You are tracking the growth of your database.
To update authorizations for users	You have 7 new users on your system, and they need authorization to perform a hexdump with !DB/WORKBENCH.
To restore an earlier extract from the DB2 catalog	You want to recover a dropped table space in an extract.
To display the system activity log	You want to resolve a series of ABENDs that occurred for a specific user.

Considerations when using !DB/WORKBENCH Housekeeping options

Review these considerations before using !DB/WORKBENCH Housekeeping options:

- Some tasks require significant time to process.
- Tasks that require significant time affect the performance of the system for other users.

Overview of !DB/WORKBENCH Housekeeping

Overview

This unit surveys the !DB/WORKBENCH Housekeeping options available for displaying information, customizing your system, using extracts, and restoring catalog copies. It provides information about the tasks for the options.

Authorization requirements

Many of the options in Housekeeping require !DB/WORKBENCH authorization. When the option that you want requires authorization, any of these !DB/WORKBENCH authorizations (or other appropriate authorizations) are sufficient:

- ECP authorization to run an extract
- ECX authorization to run an extract
- R/O (Read/Only) access to the profile dataset or to the system PDS (partitioned dataset)
- R/W (Read/Write) access to the profile dataset or to the system PDS

Options available for displaying information

This chart lists the options available for displaying information.

Task You Want to Perform	Option to Select
Display global dataset information	Option 1 (Profile Data Administration), then Option 0 (Global Dataset Information)
Display catalog table prefixes for IDF	Option 1 (Profile Data Administration), then Option 4 (Catalog Prefixes for IDF)
Display activity log	Option 7 (Display Activity Log)
Display extract history	Option 3 (Display Extract History)

Options available for customizing your system

This chart lists the options available for customizing your system.

Task You Want to Perform	Option to Select
Specify information for a specific DB2 subsystem	Option 1 (Profile Data Administration), then Option 3 (DB2-Subsystem Configuration Information)
Configure !DB/WORKBENCH	Option 1 (Profile Data Administration), then Option 5 (!DB/WORKBENCH Configuration Information)
Configure !DB/WORKBENCH processing information	Option 1 (Profile Data Administration), then Option 6 (!DB/WORKBENCH Processing Information)
Update authorizations for users	Option 2 (Update Authorizations Information)
Manipulate data in the activity log	Option 7 (Display Activity Log)

Options available for using extracts

This chart lists the options available for using extracts. It includes information about restoring catalog copies.

Task You Want to Perform	Option to Select
Display historical information about an earlier extract	Option 3 (Display Extract History)
Run a new DB2 catalog extract and extract referential integrity information	Option 4 (Run a New Extract)
Restore an earlier DB2 catalog extract	Option 6 (Resurrect old copy of the Catalog)

Accessing !DB/WORKBENCH Housekeeping Options

Overview

This unit contains information about accessing the options for !DB/WORKBENCH Housekeeping. It includes information on accessing the menu for !DB/WORKBENCH Housekeeping options and methods for accessing the Housekeeping options.

Information displayed with the !DB/WORKBENCH Housekeeping options menu

The !DB/WORKBENCH Housekeeping options menu has housekeeping options that are listed by option number. You select the option you want from the menu.

Accessing the !DB/WORKBENCH Housekeeping options menu

Follow these steps to access the Housekeeping options menu from the !DB/WORKBENCH Primary Menu.

Step	Action
1	On the !DB/WORKBENCH Primary Menu, type H in the Option field.
2	<p>Press Enter.</p> <p>Result: The Housekeeping options menu is displayed.</p> <div data-bbox="573 1209 1252 1556" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> ADMINISTRATIVE FUNCTIONS-- DB/WORKBENCH --- DB2=DB41A----- Cmd ==> Option ==> 1 - Profile Data Administration 2 - Update the Authorizations Information 3 - Display Extract History - The last extract was run on 07/16/99 at 2:14 4 - Run a new Extract 5 - Display System PDS Members 6 - Resurrect old copy of the Catalog (LAZARUS) 7 - Display the Activity Log of the LOG PDS 8 - Display DB/Tools profile data sets 9 - Set DB/Tools profile data sets </pre> </div>

Methods for accessing !DB/WORKBENCH Housekeeping options

You can access !DB/WORKBENCH Housekeeping options in these two ways:

- accessing from the !DB/WORKBENCH Housekeeping Options menu
- accessing from the !DB/WORKBENCH Primary Menu (quick access)

Before you can access a !DB/WORKBENCH Housekeeping option from the !DB/WORKBENCH Primary Menu, you must know the character of the selection you want.

Accessing options from the !DB/WORKBENCH Housekeeping Options menu

Follow these steps to access the Housekeeping options from the !DB/WORKBENCH Housekeeping Options menu.

Step	Action
1	On the !DB/WORKBENCH Housekeeping options menu, type the character of the selection you want in the Option field. Example: Type 1 to access the Profile Data Administration panel and its options.
2	Press Enter. Result: The system displays the panel or information for the selection you entered.

Accessing options from the !DB/WORKBENCH Primary Menu (quick access)

Follow these steps to quickly access the Housekeeping options from the !DB/WORKBENCH Primary Menu.

Step	Action
1	On the !DB/WORKBENCH Primary Menu, type H followed by a period and the character of the selection you want in the Option field. Example: Type H.1 to access the Profile Data Administration panel and its options.
2	Press Enter. Result: The system displays the panel or information for the selection you entered.

Option 1: Profile Data Administration

Overview

This unit contains introductory information on the !DB/WORKBENCH Housekeeping option used to display global profile information for the production !DB/Tools datasets. It also displays the DB2 subsystem and !DB/WORKBENCH configuration information.

Background about profile data administration

This option (referred to as Housekeeping Option 1) displays a list of options that are used to configure the selected DB2 subsystem and !DB/WORKBENCH. You can use these options to verify the settings or change them to customize your installation.

Displaying the profile data administration information

You can display !DB/Tools profile data administration information using this option. This information consists of:

- !DB/Tools global dataset information (option 0)
- !DB/Tools global configuration information (option 1)
- !DB/Tools global print options (option 2)
- DB2 subsystem configuration information (option 3)
- DB2 subsystem catalog prefixes for the Interactive Display Facility (IDF) (option 4)
- !DB/WORKBENCH configuration information (option 5)
- !DB/WORKBENCH processing information (option 6)

Accessing the Profile Data Administration panel

Follow these steps to access the Profile Data Administration panel (and its options) from the !DB/WORKBENCH Housekeeping options menu.

Step	Action
1	On the Housekeeping options menu, type 1 in the Option field.
2	<p>Press Enter.</p> <p>Result: The Profile Data Administration panel is displayed.</p> <div data-bbox="570 606 1252 905" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> ADMINISTRATIVE FUNCTIONS-- DB/WORKBENCH --- DB2=DB41A----- Cmd ==> Option ==> 0 Global - Data Set Information 1 Global - Configuration Information 2 Global - Print Options 3 DB2-Subsystem - Configuration Information 4 DB2-Subsystem - Catalog Prefixes for IDF 5 DB/Workbench - Configuration Information 6 DB/Workbench - Processing Information </pre> </div>

Profile Data Administration panel options

The sections immediately following this unit explain options 0–6 on the Profile Data Administration panel.

Option 1.0: Global Dataset Information

Overview

This unit describes the !DB/Tools global dataset information.

Background about global dataset information

This option (referred to as Housekeeping Option 1.0) displays the names of the production !DB/Tools datasets. The datasets contain global !DB/Tools product data for your current MVS system. Changing this data may require a change to the DB2 subsystem and profile data to ensure concurrency; therefore, be careful when making any changes to the data.

Changing the global dataset information

You can change the !DB/Tools global dataset information using this option by typing over the existing values for the field with values you provide. You can change the following information:

- names of the CLIST, CNTL, LOAD, and DBRM libraries
 - names of the production !DB/Tools datasets
 - symbolic name for virtual I/O
-

Authority required for global dataset information

To use this option for the global dataset information, you must have these authorizations:

- To display information only, you must have at least R/O access to the profile dataset.
- To change values, you must have R/W access to the profile dataset.

Accessing the !DB/Tools Global Dataset Names panel

Follow these steps to access the !DB/Tools Global Dataset Names panel from the !DB/WORKBENCH Profile Data Administration panel.

Step	Action
1	On the Profile Data Administration panel, type 0 in the Option field.
2	<p>Press Enter.</p> <p>Result: The !DB/Tools Global Dataset Names panel is displayed.</p> <div data-bbox="570 674 1248 1018" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> ----- DB/Tools ----- Cmd ==> DB/Tools Global Data Set Names CAUTION: These are the names of your production DB/Tools data sets. Please be careful when changing them. CLIST Data Set ==> TDKT.DBTOOLS.V500.CLIST CNTL Data Set ==> TDKT.DBTOOLS.V500.CNTL LOAD Data Set ==> TDKT.DBTOOLS.V500.LOAD DBRM Data Set ==> TDKT.DBTOOLS.V500.DBRM MSGs Data Set ==> TDKT.DBTOOLS.V500.MSGS PANELS Data Set ==> TDKT.DBTOOLS.V500.PANELS TABLES Data Set ==> TDKT.DBTOOLS.V500.CNTL SKELS Data Set ==> TDKT.DBTOOLS.V500.SKELS UTIL Data Set ==> TDKT.DBTOOLS.V500.UTIL </pre> </div>

Field definitions for the Global Dataset Information panel

You can change the values for the fields displayed on the panel.

Field Name	Description
CLIST Library	Name of library for product CLISTs
CNTL Dataset	Name of dataset for product JCL
LOAD Library	Name of library for !DB/Tools LOAD modules
DBRM Library	Name of library for DBRMs required to BIND the product plans
MSG5 Dataset	Name of dataset for product ISPF messages
PANELS Dataset	Name of dataset for product panels
TABLES Dataset	Name of dataset for product ISPF command and data tables
SKELS Dataset	Name of dataset for product ISPF skeleton
UTIL Dataset	Name of dataset for product utility profiles

Option 1.1: Global Configuration Information

Overview

This unit describes the !DB/Tools Global Configuration Information panel.

Background about the global configuration information

This option (referred to as Housekeeping Option 1.1) allows you to change the !DB/Tools global configuration information for your current MVS system.

Changing the global configuration information

You can change the global configuration information by typing over the existing values for the fields with values you provide.

Authority required for global configuration information

To use this option for the global configuration information, you must have these authorizations:

- To display information only, you must have at least R/O access to the profile dataset.
- To change values, you must have R/W access to the profile dataset.

Accessing the !DB/Tools Global Configuration Information panel

Follow these steps to access the Profile Data Administration panel (and its options) from the !DB/WORKBENCH Housekeeping options menu.

Step	Action
1	On the Profile Data Administration panel, type 1 in the Option field.
2	<p>Press Enter.</p> <p>Result: The !DB/Tools Global Configuration Information panel is displayed.</p> <div data-bbox="570 674 1250 970" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> ----- DB/Tools ----- Cmd ==> DB/Tools Global Configuration Information PERMANENT? ==> Y (Y Yes N No) VIO Symbolic Name ==> VIO Temp Disk Symbolic Name ==> SYSDA Max BLKSIZE for VIO ==> 0 SYSOUT Hold Class ==> X Enable Edit Recovery? ==> N (Y Yes N No) RETURN = EXIT? ==> N (Y Return exits N Return goes to Primary Menu) </pre> </div>

Field definitions for the !DB/Tools Global Configuration Information panel

You can change the values for the fields displayed on the panel.

Field Name	Description
PERMANENT?	Indicates if the changes are permanent or for the current session only.
VIO Symbolic Name	Symbolic name for virtual I/O; is specific for the current MVS system
Temp Disk Symbolic Name	Symbolic name for temporary disk files; specific to a given MVS system
Max BLKSIZE for VIO	Maximum block size to use when allocating virtual I/O datasets
SYSOUT hold class	Your installation's hold class; default is class X
Enable Edit Recovery?	Indicates whether Edit Recovery should be enabled.
RETURN = EXIT?	Indicates whether '=x' exits or goes to the Primary Menu.

Option 1.2: Global Print Options

Overview

This unit describes the !DB/Tools Global Print Options panel.

Background about the global print options

This option (referred to as Housekeeping Option 1.2) allows you to change and configure the print options for all !DB/Tools products.

Changing the global print options

You can change the !DB/Tools global print options by typing over the existing values with values you provide.

Authority required for global print options

To use this option for the global print options, you must have these authorizations:

- To display information only, you must have at least R/O access to the profile dataset.
- To change values, you must have R/W access to the profile dataset.

Accessing the !DB/Tools Global Print Options panel

Follow these steps to access the !DB/Tools Global Print Options panel from the !DB/WORKBENCH Housekeeping options menu.

Step	Action
1	On the Profile Data Administration panel, type 2 in the Option field.
2	<p>Press Enter.</p> <p>Result: The !DB/Tools Global Print Options panel is displayed.</p> <div data-bbox="573 642 1252 989" style="border: 1px solid black; padding: 10px;"> <pre> ----- DB/Tools ----- Cmd ==> DB/Tools Global Print Options PERMANENT? ==> N (Y Yes N No) Print Class ==> X Class (A-Z, 0-9) Lines per page ==> 60 Lines per page Special Forms ID ==> Output Form number Destination ID ==> Printer name or Remote Station Print to User ID ==> Route to User ID Writer ==> Writer name Allow User ID and ==> N (Y Yes N No) Destination ID? </pre> </div>

Field definitions for the !DB/Tools Global Print Options panel

You can change the values for the fields displayed on the panel.

Field Name	Description
PERMANENT?	Indicates if the changes are permanent or for the current session only.
Print Class	The JES output class where the output from the PRNT command is to be placed.
Lines per Page	The number of lines per page for printed output.
Special Forms ID	The special forms to use for the output from the PRNT command. This may be any value allowed by your installation for JES output.
Destination ID	The JES printer ID or destination for printed output. This may be any value allowed by your installation for JES output.
Print to User ID	The printer destination user ID for printed output. This may be any value allowed by your installation for JES output.
Writer	The JES writer name. This may be any value allowed by your installation for JES output.
Allow User ID and Destination ID?	Indicates whether to allow both the User ID and Destination ID to be input by the user on PRNT commands.

Option 1.3: DB2 Subsystem-Configuration Information

Overview

This unit contains introductory information on the !DB/WORKBENCH Housekeeping option for changing information for profile data items specific to a single DB2 subsystem.

Background about the specific DB2 subsystem

This option (referred to as Housekeeping Option 1.3) allows you to change field values for profile data pertaining to a single DB2 subsystem (for example, extract ID). All DB2 configuration information is based upon the extract ID you select. Values for the fields are required for operation of !DB/Tools. Unqualified changes of the values may affect the integrity of your system; therefore, be careful when changing values.

Changing information for the specific DB2 subsystem

You can change DB2 subsystem information by typing over the existing values for the fields with values you provide. You can change the IDs for the DB2 subsystem and the MVS subsystem and the names of the datasets that contain DSN1COPY, DSNTIAR, and the DSNZPARMs members.

You can change these values:

- name of the DB2 subsystem location, of the DSNZPARM member, and of the database that contains !DB/Tools objects
 - version of DB2 you want to run
 - prefix of the DB2 catalog
-

Authority required for the specific DB2 subsystem

To use the option for specific DB2 subsystem information, you must have these authorizations:

- To display information only, you must have at least R/O access to the profile dataset.
- To change values, you must have R/W access to the profile dataset.

Accessing the !DB/Tools DB2 Configuration Information panel

Follow these steps to access the !DB/Tools DB2 Configuration Information panel from the !DB/WORKBENCH Profile Data Administration panel.

Step	Action
1	On the Profile Data Administration panel, type 3 in the Option field.
2	<p>Press Enter.</p> <p>Result: The !DB/Tools DB2 Configuration Information panel is displayed.</p> <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> ----- DB/Tools ----- Cmd ==> DB/Tools DB2 Configuration Information Extract ID ==> DB41 MVS ID ==> SYSA Extract Description ==> SYSA DB2 Subsystem ID ==> D41A DB2 Version ==> 410 DB2 Location Name ==> D41A DSNTIAUL Plan Name ==> DSNTIB41 Real Catalog Prefix ==> SYSIBM DSNTIAD Plan Name ==> DSNTIA41 Catalog Prefix ==> SYSIBM DSNZPARM Member Name ==> DSNZPARM DB/Tools Database ==> V5000BDB DB/Tools Qualifier ==> TUSER Collection Id Prefix ==> V500 DB2 DSNLOAD Data Set ==> PP.DB2.V4R1.SDSNLOAD DB2 DSNEXIT Data Set ==> TDD241A.DB2410.DSNEXIT DB2 RUNLIB Data Set ==> PP.DB2.V410.RUNLIB.LOAD Data Set Containing DSNHDECP ==> PP.DB2.V4R1.SDSNLOAD Data Set Containing DSN1COPY ==> PP.DB2.V4R1.SDSNLOAD Data Set Containing DSNTIAUL ==> TDD241A.DB2410.RUNLIB.LOAD Data Set Containing DSNTIAD ==> TDD241A.DB2410.RUNLIB.LOAD Data Set Containing DSNZPARM ==> TDD241A.DB2410.DSNEXIT </pre> </div>

Field definitions for the !DB/Tools DB2 Configuration Information panel

You can change the values for fields displayed on the panel, unless indicated as a display only field.

Field Name	Description
Extract ID	Identifier of the current DB2 subsystem or extract set (This field is display only and all information on the panel relates to the specific extract ID displayed.)
MVS ID	ID of the current MVS system (This ID is for local documentation only.)
Extract Description	Description of current extract ID

Field definitions for the !DB/Tools DB2 Configuration Information panel (continued)

Field Name	Description
DB2 Subsystem ID	ID of the current DB2 subsystem (This ID is used for communications with DB2.)
DB2 Version	DB2 version for the current DB2 subsystem
DB2 Location Name	Location name of the local DB2 subsystem (used for DB2 V2.3 or higher)
DSNTIAUL Plan Name	Name of the DSNTIAUL plan
Real Catalog Prefix	The prefix of the actual DB2 catalog tables if the prefix specified for Catalog Prefix is for a shadow catalog
DSNTIAD Plan Name	Name of the DSNTIAD plan
Catalog Prefix	Prefix of the DB2 catalog
DSNZPARM Member Name	Name of the member that contains DSNZPARMs
DB/Tools Database	Name of the DB2 database that contains the !DB/Tools objects
DB/Tools Qualifier	Name of the authid (used as the qualifier of the DB/Tools Tables; for example, DSPSTATS and SPCSTATS)
Collection ID Prefix	Prefix for constructing called ID for packages
DB2 DSNLOAD Dataset	Name of the dataset that contains DSNLOAD
DB2 DSNEXIT Dataset	Name of the dataset that contains DSNEXIT
DB2 RUNLIB Dataset	Name of the dataset that contains DB2 RUNLIB
Dataset Containing DSNHDECP	Name of the dataset that contains DSNHDECP
Dataset Containing DSN1COPY	Name of the dataset that contains DSN1COPY

Field definitions for the !DB/Tools DB2 Configuration Information panel (continued)

Field Name	Description
Dataset Containing DSNTIAUL	Name of the dataset that contains DSNTIAUL
Dataset Containing DSNTIAD	Name of the dataset that contains DSNTIAD
Dataset Containing DSNZPARMs	Name of the dataset that contains the DSNZPARM LOAD module (DSNZPARM is the LOAD module that contains initialization parameters for the macros that include DSN6ARVP, DSN6ENV, DSN6LOGP, DSN6SPRM, DSN6FAC, and DSN6SYSP.)

Option 1.4: Catalog Prefixes for the IDF

Overview

This unit contains introductory information on the !DB/WORKBENCH Housekeeping option for catalog prefixes for IDF (interactive display facility). It provides information about changing values for the catalog table prefixes for DB2 local and remote tables.

Background about catalog prefixes for IDF

This option (referred to as Housekeeping Option 1.4) allows you to change values for the catalog table prefixes you need to access the DB2 catalog interactively.

Changing the values for catalog prefixes for IDF

With this !DB/WORKBENCH option, you can display listings of local and remote DB2 tables and change the values for the catalog table prefixes by typing over the existing values with values you provide.

Authority required for catalog prefixes for IDF

To use the option for catalog prefixes for IDF, you must have these authorizations:

- To display information only, you must have at least R/O access to the profile dataset.
- To change values, you must have R/W access to the profile dataset.

Accessing the Catalog Prefixes for the IDF panel

Follow these steps to access the Catalog Prefixes for the IDF panel from the !DB/WORKBENCH Profile Data Administration panel.

Step	Action
1	On the Profile Data Administration panel, type 4 in the Option field.
2	<p>Press Enter.</p> <p>Result: The Catalog Prefixes for the IDF panel is displayed.</p> <div data-bbox="570 640 1248 1087" style="border: 1px solid black; padding: 10px;"> <pre> ----- DB/Tools Profile Information -- DB2 ----- Cmd ==> CATALOG PREFIXES for the Interactive Display Facility (IDF) LOCAL DB2 TABLE INFORMATION SYSCHECKS SYSIBM SYSCHECKDEP SYSIBM SYSCOLAUTH SYSIBM SYSCOLUMNS SYSIBM SYSCOPY SYSIBM SYSDATABASE SYSIBM SYSDBAUTH SYSIBM SYSDBRM SYSIBM SYSFIELDS SYSIBM SYSFORKEYS SYSIBM SYSINDEXES SYSIBM SYSINDEXPART SYSIBM SYSKEYS SYSIBM SYSLINKS SYSIBM SYSPACKAGE SYSIBM SYSPACKAUTH SYSIBM SYSPACKDEP SYSIBM SYSPACKLIST SYSIBM SYSPACKSTMT SYSIBM SYSPKSYSTEM SYSIBM SYSPLAN SYSIBM SYSPLANAUTH SYSIBM SYSPLANDEP SYSIBM SYSPLSYSTEM SYSIBM SYSPROCEDURES SYSIBM SYSRELS SYSIBM SYSRESAUTH SYSIBM SYSSTMT SYSIBM SYSSTOGROUP SYSIBM SYSSTRINGS SYSIBM SYSSYNONYMS SYSIBM SYSTABAUTH SYSIBM SYSTABLEPART SYSIBM SYSTABLES SYSIBM SYSTABLESPACE SYSIBM SYSUSERAUTH SYSIBM SYSVIEWDEP SYSIBM SYSVIEWS SYSIBM SYSVLTREE SYSIBM SYSVOLUMES SYSIBM SVSTREE SYSIBM REMOTE DB2 TABLE INFORMATION SYSLUNAMES SYSIBM SYSLOCATIONS SYSIBM SYSLUMODES SYSIBM SYSLMODESELECT SYSIBM SYSUSERS SYSIBM </pre> </div>

Values required for catalog prefixes for IDF

The values required for this option are the values for the catalog table prefixes.

Option 1.5: !DB/WORKBENCH Configuration Information

Overview

This unit contains introductory information on using the !DB/WORKBENCH Housekeeping option to specify the configuration of your system.

Background about configuration information

This option (referred to as Housekeeping Option 1.5) allows you to change values that are used to configure your system. Unqualified changes of the values may affect the integrity of your system; therefore, be careful when changing values.

Description of the process of changing configuration information

You can change configuration information by typing over the existing values in the fields on the panel. When you type the values and press Enter, the system saves the changes you have made to the profile data set.

Changing configuration information

You can change any value for the configuration information; however, refer to the *!DB/Tools Installation and Customization Guide* for information on changing the values for Access Exit, Options Exit, Extract Exit, and R-O PDS Exit.

Authority required for configuration information

To use the option for configuration information, you must have these authorizations:

- To display information only, you must have at least R/O access to the profile dataset.
 - To change values, you must have R/W access to the profile dataset.
-

Lockout conditions concerning the PDS files

It is advisable to have ample space available in your system PDS, user PDS, and log PDS files. When these operate at near capacity, you run the risk of creating a lockout condition in your TSO session for one or more users. If !DB/WORKBENCH has been used extensively, you can clean up these files to create more room, or you can increase the size of these files to prevent this condition from occurring.

Accessing the !DB/WORKBENCH Configuration Information panel

Follow these steps to access the !DB/WORKBENCH Configuration Information panel from the !DB/WORKBENCH Profile Data Administration panel.

Step	Action
1	On the Profile Data Administration panel, type 5 in the Option field.
2	<p>Press Enter.</p> <p>Result: The !DB/WORKBENCH Configuration Information panel is displayed.</p> <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> ----- DB/WORKBENCH - CONFIGURATION INFORMATION ----- Cmd ==>> Extract PLAN Name ==> KTWPLANQ Plan Name for the DB2 Extract SQL PLAN Name ==> KTWPLANQ Plan Name for the Dynamic SQL Facility System PDS ==> TDKT.DBTOOLS.W500.D41A.WKB.SYSTEM User PDS ==> TDKT.DBTOOLS.W500.D41A.WKB.USER Log PDS ==> TDKT.DBTOOLS.W500.D41A.WKB.LOG Access Exit ==> STD Access Control Exit (Users to Databases) Options Exit ==> STD Exit name for this screen Extract Exit ==> STD Exit to authorize users to extract Catalog R-0 PDS Exit ==> STD Exit to prohibit System PDS updates Passwords ==> Y Display DB2 Passwords? Y/N LAZARUS Retention ==> 4 1 - 32000 Number of extracts to retain Date Format ==> 1 1 - MM/DD/YY 2 - DD/MM/YY 3 - YY/MM/DD SQLID source ==> 2 SQLID source: 1 - Creator 2 - Created By </pre> </div>

Field definitions for the !DB/WORKBENCH Configuration Information panel

You can change the values for fields displayed on the panel, unless indicated as a display only field.

Field Name	Description on the Panel
Extract PLAN Name	DB2 plan used to run the extract
SQL PLAN Name	DB2 plan used for dynamic SQL access to DB2
System PDS	Name of the system extract partitioned dataset (fully qualified name, no quotes)
User PDS	Name of the user partitioned dataset (fully qualified name, no quotes)

Field definitions for the !DB/WORKBENCH Configuration Information panel (continued)

Field Name	Description on the Panel
Log PDS	Name of the log partitioned dataset (fully qualified name, no quotes)
Access Exit	Exit you can use to control (or authorize) user access to user databases
Options Exit	Exit you can use to control user access to this panel
Extract Exit	Exit you can use to control (or authorize) user access to catalog extract
R-O PDS Exit	Exit to prohibit System PDS updates
Passwords	Display DB2 Passwords? Y/N
Lazarus Retention	1–32000 Number of Lazarus generations to save
Date Format	1 - mm/dd/yy 2 - dd/mm/yy 3 - yy/mm/dd The option number sets the format for the date on panels with a date field.
SQLID source	1 - Creator 2 - Created by

Background about the Lazarus Retention field and the STATDBAS table

The value in the Lazarus Retention field determines the number of generations of the extract that !DB/WORKBENCH stores in the system PDS. The value has no effect on reduction of data that !DB/WORKBENCH inserts in the tables in the STATDBAS table space.

For information about controlling the insertion of extract data in the tables in the STATDBAS table space, see the *!DB/Tools Installation and Customization Guide*. For general information about the STATDBAS table space or about reducing data, see “Managing Statistics in the STATDBAS Table Space” on page 195.

Option 1.6: !DB/WORKBENCH Processing Information

Overview

This unit contains introductory information on using the !DB/WORKBENCH Housekeeping option which allows you to communicate special processing information to !DB/WORKBENCH.

Background about processing information

This option (referred to as Housekeeping Option 1.6) allows you to set values for special processing information in !DB/WORKBENCH, allowing you to further customize the product.

Changing processing information

You can change values for the processing information by typing over the existing values with values you provide; however, refer to the *!DB/Tools Installation and Customization Guide* for information on changing the values for Access Exit, Options Exit, Extract Exit, and R-O PDS Exit.

Accessing the !DB/WORKBENCH Processing Information panel

Follow these steps to access the !DB/WORKBENCH Processing Information panel from the !DB/WORKBENCH Profile Data Administration panel.

Step	Action
1	On the Profile Data Administration panel, type 6 in the Option field.
2	<p>Press Enter.</p> <p>Result: The !DB/WORKBENCH Processing Information panel is displayed.</p> <div data-bbox="571 674 1250 1121" style="border: 1px solid black; padding: 10px;"> <pre> ----- DB/WORKBENCH - PROCESSING INFORMATION ----- Cmd ==> Hiperspace Control ==> N Y/N to bypass use of Hiperspace for the Colume display. Grants ==> N Y/N to generate Grants when the GRANTOR = GRANTEE (implicit Grants) during OUTDDL processing. Package Authority ==> N Y/N to include packages in the standard auth. processing. WARNING - this could seriously degrade processing performance. Batch Return Code ==> N Y/N to always set the batch processing return code to 0. N will allow standard return codes to be generated. Temp. Work Unit ==> N Y/N to use SYSDA rather than VIO for a temporary work unit (device). RACF Qualifier ==> N Y/N to use a RACF Qualifier (up to 8 Characters) QUAL ==> Utility Job Card ==> N Y/N to retain the Job Card from the previous ISPF session for the utilities job. Extract Sort ==> N Y/N to use a DB2 sort for the extract. Y will use an external sort for the extract. Show Line Mode ==> N Y/N to display line mode error messages. Y will display sysprint type error messages. </pre> </div>

Field definitions for the !DB/WORKBENCH Processing Information panel

You can change values for fields displayed on the panel. The default value for each field is N (no) or blank.

Field Name	Description on the Panel
Hiperspace Control	Type Y to bypass use of Hiperspace
Grants	Type Y to generate grants when Grantor=Grantee
Package Authority	Type Y to perform authorization checks on packages
Batch Return Code	Type Y to always set batch return codes to zero
Temp. Work Unit	Type Y to use SYSDA instead of virtual I/O (VIO)
RACF Qualifier	Type Y to use RACF qualifier (If you specify Y, you must enter the name, up to 8 characters, of the RACF qualifier in the QUAL field.)
Utility Job Card	Type Y to use the job card from the last session for utilities
Extract Sort	Type Y to use an external sort
Show Line Mode	Type Y to display line mode error messages

Option 2: Update Authorizations Information

Overview

This unit refers to the option that lets you control user access values in !DB/WORKBENCH.

Information for using this option

For instructions on using this option, see “Appendix D. !DB/Tools Access Control Statements” in the *!DB/Tools Installation and Customization Guide*.

Option 3: Display Extract History

Overview

This unit contains introductory information about using !DB/WORKBENCH to display historical information and extracts of the DB2 catalog. It provides prerequisite information you will need to use this option. This unit also includes information about extracts and the system PDS.

Background about displaying extract history

This option for displaying historical information about catalog extracts allows you to select an extract and to display detailed historical information and the extract. The information displayed on the panel for this option contains extracts of the catalog listed by the dates the extracts were run. The following objects are displayed in each extract:

- databases
 - table spaces
 - tables
 - indexes
 - keys
 - datasets (The approximate number is displayed.)
-

Background about displaying data with !DB/WORKBENCH

!DB/WORKBENCH displays more data than is available if you displayed extract history using SPUFI. It includes in its totals both the number of objects extracted from the DB2 catalog (as does SPUFI) and also the number of objects in the system database DSNDB01.

Authority required for displaying extract history

To use the option for displaying historical information, you must have at least R/O access to the system PDS.

Accessing the Display Extract History panel

Follow these steps to access the Extract History panel from the !DB/WORKBENCH Housekeeping options menu.

Step	Action																																																																																																																																								
1	On the Housekeeping options menu, type 3 in the Option field.																																																																																																																																								
2	<p data-bbox="570 489 727 520">Press Enter.</p> <p data-bbox="570 537 1198 573">Result: The Extract History panel is displayed.</p> <div data-bbox="570 604 1247 1066" style="border: 1px solid black; padding: 5px;"> <pre data-bbox="597 632 1219 667">EXTRACT HISTORY ----- DB/WORKBENCH --- DB2=D41A----- EXTRACT 1 OF 43 Cmd ==>> Scroll ==>> PAGE</pre> <p data-bbox="597 682 743 703">Sels: S Details</p> <table border="1" data-bbox="597 716 1219 1045"> <thead> <tr> <th>SEL</th> <th>DATE</th> <th>NO. OF DATABASES</th> <th>NO. OF TBLSPCS</th> <th>NO. OF TABLES</th> <th>NO. OF INDEXES</th> <th>NO. OF KEYS</th> <th>APPROX # DATASETS</th> </tr> </thead> <tbody> <tr><td>-</td><td>10/16/99</td><td>79</td><td>341</td><td>498</td><td>516</td><td>2,327</td><td>857</td></tr> <tr><td>-</td><td>08/19/99</td><td>58</td><td>197</td><td>392</td><td>430</td><td>1,871</td><td>627</td></tr> <tr><td>-</td><td>08/18/99</td><td>58</td><td>197</td><td>392</td><td>430</td><td>1,871</td><td>627</td></tr> <tr><td>-</td><td>08/17/99</td><td>58</td><td>197</td><td>392</td><td>430</td><td>1,871</td><td>627</td></tr> <tr><td>-</td><td>08/16/99</td><td>51</td><td>173</td><td>377</td><td>408</td><td>1,852</td><td>581</td></tr> <tr><td>-</td><td>08/09/99</td><td>50</td><td>173</td><td>377</td><td>408</td><td>1,852</td><td>581</td></tr> <tr><td>-</td><td>08/04/99</td><td>49</td><td>173</td><td>377</td><td>408</td><td>1,855</td><td>581</td></tr> <tr><td>-</td><td>08/03/99</td><td>49</td><td>173</td><td>377</td><td>406</td><td>1,844</td><td>579</td></tr> <tr><td>-</td><td>08/01/99</td><td>49</td><td>172</td><td>375</td><td>404</td><td>1,842</td><td>576</td></tr> <tr><td>-</td><td>07/31/99</td><td>49</td><td>172</td><td>379</td><td>404</td><td>1,842</td><td>576</td></tr> <tr><td>-</td><td>07/19/99</td><td>48</td><td>171</td><td>379</td><td>404</td><td>1,842</td><td>575</td></tr> <tr><td>-</td><td>07/13/99</td><td>46</td><td>167</td><td>363</td><td>398</td><td>1,820</td><td>565</td></tr> <tr><td>-</td><td>05/01/99</td><td>45</td><td>161</td><td>320</td><td>355</td><td>1,491</td><td>516</td></tr> <tr><td>-</td><td>04/28/99</td><td>45</td><td>161</td><td>320</td><td>355</td><td>1,491</td><td>516</td></tr> <tr><td>-</td><td>04/27/99</td><td>45</td><td>161</td><td>320</td><td>355</td><td>1,491</td><td>516</td></tr> <tr><td>-</td><td>04/25/99</td><td>45</td><td>161</td><td>320</td><td>355</td><td>1,491</td><td>516</td></tr> </tbody> </table> </div>	SEL	DATE	NO. OF DATABASES	NO. OF TBLSPCS	NO. OF TABLES	NO. OF INDEXES	NO. OF KEYS	APPROX # DATASETS	-	10/16/99	79	341	498	516	2,327	857	-	08/19/99	58	197	392	430	1,871	627	-	08/18/99	58	197	392	430	1,871	627	-	08/17/99	58	197	392	430	1,871	627	-	08/16/99	51	173	377	408	1,852	581	-	08/09/99	50	173	377	408	1,852	581	-	08/04/99	49	173	377	408	1,855	581	-	08/03/99	49	173	377	406	1,844	579	-	08/01/99	49	172	375	404	1,842	576	-	07/31/99	49	172	379	404	1,842	576	-	07/19/99	48	171	379	404	1,842	575	-	07/13/99	46	167	363	398	1,820	565	-	05/01/99	45	161	320	355	1,491	516	-	04/28/99	45	161	320	355	1,491	516	-	04/27/99	45	161	320	355	1,491	516	-	04/25/99	45	161	320	355	1,491	516
SEL	DATE	NO. OF DATABASES	NO. OF TBLSPCS	NO. OF TABLES	NO. OF INDEXES	NO. OF KEYS	APPROX # DATASETS																																																																																																																																		
-	10/16/99	79	341	498	516	2,327	857																																																																																																																																		
-	08/19/99	58	197	392	430	1,871	627																																																																																																																																		
-	08/18/99	58	197	392	430	1,871	627																																																																																																																																		
-	08/17/99	58	197	392	430	1,871	627																																																																																																																																		
-	08/16/99	51	173	377	408	1,852	581																																																																																																																																		
-	08/09/99	50	173	377	408	1,852	581																																																																																																																																		
-	08/04/99	49	173	377	408	1,855	581																																																																																																																																		
-	08/03/99	49	173	377	406	1,844	579																																																																																																																																		
-	08/01/99	49	172	375	404	1,842	576																																																																																																																																		
-	07/31/99	49	172	379	404	1,842	576																																																																																																																																		
-	07/19/99	48	171	379	404	1,842	575																																																																																																																																		
-	07/13/99	46	167	363	398	1,820	565																																																																																																																																		
-	05/01/99	45	161	320	355	1,491	516																																																																																																																																		
-	04/28/99	45	161	320	355	1,491	516																																																																																																																																		
-	04/27/99	45	161	320	355	1,491	516																																																																																																																																		
-	04/25/99	45	161	320	355	1,491	516																																																																																																																																		

Displaying detailed historical information about an extract

Follow the steps to display detailed historical information about the extract you select.

Step	Action
1	Type S in the select field next to the extract.
2	<p>Press Enter.</p> <p>Result: The detailed historical information is displayed for the extract you selected. This is a display-only panel.</p> <div data-bbox="573 642 1248 984" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> EXTRACT HISTORY DETAILS----- DB/WORKBENCH --- DB2=D41A----- ROW 1 OF 14 Cmd ==> Scroll ==> PAGE EXTRACT DATE: 08/19/99 DATABASES: 58 APPROX. DATASETS: 627 TABLESPACES: 197 CARDS IN MEMBER: 10,020 TABLES: 392 COLUMNS: 18,976 INDEXES: 430 VIEWS: 45 PLANS: 611 SYNONYMS: 27 PACKAGES: 302 SYSFIELDS: 50 ALIASES: 31 KEYS: 1,871 STORAGE GROUPS: 29 </pre> </div>

Background about extract information in the system PDS

When you run a new extract, information about the extract, such as the timestamp, is saved in members in the !DB/WORKBENCH system PDS. The system PDS also contains information about Lazarus copies.

Elements of the system PDS that contain extract information

The illustration shows the elements of the system PDS that contain extract information.

```

EDIT --- TDKT.DBTOOLS.V31.DB2.WKB.SYSTEM ----- ROW 004690F 00577
COMMAND ==>>                                SCROLL ==>> CSR
NAME          VV.MM  CREATED    CHANGED    SIZE  INIT  MOD  ID

.<..D...
.(..D...
+..D...
|..D... 1
&..D...
&..D...
...D...
.
.
.
.
2 .m...DATB
.
.m. | ..DD 3
.m. | ..DF
.m. | ..DU
.m. | ..TB
.
.
.m.?..DD 4
.m.?..DF
.m.?..DU
.m.?..TB
.
.
.
.DB2COLS
.DB2DATB
.DB2DATF
.DB2DBLD 5
.DB2DBLF
.DB2DBLU
.DB2OLDD
.
.
.
.Z..... 6
**END**
    
```

Elements of the system PDS that contain extract information (continued)

This list defines the elements of the system PDS that contain extract information.

- 1** PDS members for a database (There is one PDS member for each database in the system.)
- 2** Unprintable characters (such as .m...) for the timestamp
- 3** Set of four PDS members for a Lazarus copy
- 4** Characters (such as .DB2) indicating it is a member for the current extract
- 5** Set of PDS members for the current extract
- 6** Member that provides a summary listing of all databases

Option 4: Run a New Extract

Overview

This unit contains introductory information on the !DB/WORKBENCH Housekeeping option for running a new extract from the DB2 catalog. It provides background and prerequisite information and the commands you will need for using the option.

Background about running a new extract

You can run a new extract from the DB2 catalog using an SQL program that selects the data from tables in the DB2 catalog. The SQL program that extracts the data may take a long time to complete the extract. Do not run the extract during prime time for usage of your system unless it is necessary.

Conditions requiring a new extract

If you run any compression routine against the !DB/WORKBENCH System PDS, you must run a new extract. Failure to do this can result in an assortment of display irregularities showing distorted data.

Background about extracting referential integrity information

When you run a new extract, referential integrity information is included as part of the !DB/WORKBENCH catalog extract. You need not run a separate extract to get referential integrity information.

Implication of the timestamp on number of extracts

When you run an extract for a DB2 subsystem, !DB/WORKBENCH creates the required members for the extract and identifies the members with a timestamp. The timestamp shows the date. Since the timestamp records the day and time to the minute, the number of extracts that you can maintain is practically unlimited.

Authority required for running a new extract

To use the option for running a new extract, you must have SELECT authority to access the DB2 catalog tables. If you do not obtain SELECT authority and you use this option, the processing for the extract will fail.

Commands to use for running a new extract

You can use the Enter key or the END command to run a new extract at this time or later. Review the following chart to determine the appropriate action to take.

Task You Want to Perform	Action to Take
Run a new extract now	Press Enter
Run a new extract later	Use the END command

Background about gathering statistical data during an extract

You define information about the extract process when you install !DB/WORKBENCH. One option you can define is whether to collect statistics as a byproduct of the extract. If you choose not to collect statistics, you will receive an informational return code of 4 from the extract process.

If you want information about the statistical data that the system can insert in the tables in the STATDBAS table space when you run an extract, see “Managing Statistics in the STATDBAS Table Space” on page 195.

Option 5: Display System PDS Members

Overview

This unit contains introductory information about displaying the members in the !DB/WORKBENCH system partitioned dataset (PDS).

Background about displaying system PDS members

The option for displaying system partitioned dataset (PDS) members allows you to manipulate the groups of members in the system partitioned dataset (PDS). The members are grouped by database. The database name is displayed in hexadecimal values.

Authority required for displaying the system PDS

To use the option for displaying the system PDS, you must have at least R/O access to the system PDS.

Accessing the System dataset Information panel

Follow these steps to access the System Dataset Information panel from the !DB/WORKBENCH Housekeeping options menu.

Step	Action
1	On the Housekeeping options menu, type 5 in the Option field.
2	<p>Press Enter.</p> <p>Result: The System Dataset Information panel is displayed.</p> <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> SYSTEM DATA SET INFORMATION----- ROW 1 OF 294 Cmd ==> Scroll ==> PAGE SYSTEM (EXTRACT) PDS NAME: TDKT.DBTOOLS.W500.D41A.WKB.SYSTEM The members of the System (Extract) PDS are grouped into logical members. S to display members in the group (size is not displayed). R to display members in the group (size is displayed). E to delete all the members of a group. SEL MEMBER NAME MBR COUNT TEXT ----- - 00010000C4000000 1 DBD = AAAAMMM - 00020000C4000000 1 DBD = A500 - 00040000C4000000 1 DBD = CANDLE06 - 00050000C4000000 1 DBD = CANDLE07 - 00060000C4000000 1 DBD = CANU500 - 00070000C4000000 1 DBD = CANV500 - 00080000C4000000 1 DBD = CBDDDB - 00090000C4000000 1 DBD = CDORMDB0 - 000A0000C4000000 1 DBD = DBGLEN - 000B0000C4000000 1 DBD = DBKARE4 - 000C0000C4000000 1 DBD = DBKENT - 000D0000C4000000 1 DBD = DBKENT2 </pre> </div>

Methods for manipulating the system PDS members

Review the chart for methods to manipulate the system PDS members.

Task You Want to Perform	Method	Name
Display a group of members and their sizes	Select	R
Display a group of members only	Select	S
Erase all members of a particular group	Select	E

Option 6: Resurrect Old Catalog Copies

Overview

This unit contains introductory information on resurrecting and erasing old copies of the !DB/WORKBENCH catalog extract. It also includes information on using the current system PDS or a different system PDS to restore catalog copies.

Background about resurrecting old catalog copies

The option for resurrecting old catalog copies allows you to manipulate the copies of the !DB/WORKBENCH catalog extract that have been saved in compressed format. You can recover or erase any of the extracts displayed on the list.

To access the latest catalog extract, you must exit !DB/WORKBENCH and then access !DB/WORKBENCH.

Authority required for restoring catalog copies

To use the option for restoring catalog copies, you must have R/W access to the system PDS.

Prerequisite for resurrecting old catalog copies

The format for of the !DB/WORKBENCH extract changed to support larger DB2 subsystems with the CUM 9808-0 tape. If you want to use Lazarus data that was saved prior to this change, you must first convert this data. See the *!DB/Tools Installation and Customization Guide* for instructions.

Accessing the (LAZARUS) Catalog Data panel

Follow these steps to access the (LAZARUS) Catalog Data panel from the !DB/WORKBENCH Housekeeping options menu.

Step	Action
1	On the Housekeeping options menu, type 6 in the Option field.
2	<p>Press Enter.</p> <p>Result: The (LAZARUS) Catalog Data panel is displayed.</p> <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> CATALOG DATA ----- DB/WORKBENCH --- DB2=D41A----- DATE 1 OF 4 Cmd ==> Scroll ==> PAGE The extracted DB2 Catalog has been saved (in a compressed format) for the dates listed below. To access a previous extract place an S beside the extract required or to erase a previous extract place an E beside the extract to erase. Press ENTER after you have made your selection. Press END without selecting any extract to have the latest Catalog Extract displayed. SEL DATE TIME --- - - 99.228 08/16/99 - 99.229 08/17/99 - 99.230 08/18/99 - 99.231 06:03 08/19/99 ***** BOTTOM OF LIST ***** </pre> </div>

Methods for manipulating old copies of the !DB/WORKBENCH extract

Review the chart for methods to manipulate old copies of the !DB/WORKBENCH extract.

Task You Want to Perform	Method	Name
Resurrect a copy of the !DB/WORKBENCH catalog extract	Select	S
Erase a copy of the !DB/WORKBENCH catalog extract	Select	E

Additional information about using the current system PDS or a different system PDS

With !DB/WORKBENCH, you can restore catalog copies using a system PDS that you choose. You can restore copies using one of these sources:

- A different system PDS
- The current system PDS

If you want to use a different system PDS to restore copies, you provide the name for the PDS.

Background about using a different system PDS

If you want to use a different system PDS to restore copies, follow all of the instructions in this unit.

Authority required for using a different system PDS

To use a different system PDS, you provide the name for a system PDS. For information about the name for a system PDS, see “Option 1.5: !DB/WORKBENCH Configuration Information” on page 165.

Procedure for restoring catalog copies using a different system PDS

Follow these steps to restore a catalog copy using a different system PDS.

Step	Action
1	On the !DB/WORKBENCH Configuration Information panel, type the name of the system PDS in the System PDS field.
2	Press Enter. Result: The system displays the !DB/WORKBENCH Primary Menu. It also displays a message indicating that the changes you made are saved.
3	On the !DB/WORKBENCH Primary Menu, type H.6 in the Option field.
4	Press Enter. Result: The system displays the panel for resurrecting old catalog copies.
5	Follow the instructions in this unit for restoring catalog copies using the current system PDS.

Procedure for restoring catalog copies using the current system PDS

Follow these steps to restore a catalog copy using the current system PDS.

Step	Action
1	Type S in the Select field for the catalog copy you want to restore.
2	Press Enter. Result: The system displays the !DB/WORKBENCH Primary Menu. It also displays a message indicating that the extract ID for the catalog copy is initialized.

Option 7: Display the Activity Log

Overview

This unit contains introductory information about displaying the activity log in the !DB/WORKBENCH log PDS (partitioned dataset).

Background about displaying the activity log

The option for displaying the activity log allows you to manipulate a list of actions that updated the DB2 system. This includes CREATEs, BINDs, AMS DEFINEs and DELETEs, DROPs, FREEs, GRANTs, and Utility Submits. Product ABENDs are also listed here. Each action has a corresponding member in the !DB/WORKBENCH log PDS.

Background about the activity log members

When you update an activity log member in the !DB/WORKBENCH log PDS, !DB/WORKBENCH saves the changes you make in the *useridP* member in the !DB/WORKBENCH user PDS.

If you want to execute the statements for the activity log member you updated, use the *useridP* member in the !DB/WORKBENCH user PDS.

Authority required for displaying the activity log

To use the option for displaying the activity log, you must have these authorizations:

- To display information only, you must have at least R/O access to the log PDS.
- To change values, you must have R/W access to the log PDS.

Accessing the Log Activity List panel

Follow these steps to access the Log Activity List panel from the !DB/WORKBENCH Housekeeping options menu.

Step	Action
1	On the Housekeeping options menu, type 7 in the Option field.
2	<p>Press Enter.</p> <p>Result: The Log Activity List panel is displayed.</p> <div data-bbox="573 611 1250 1056" style="border: 1px solid black; padding: 10px;"> <pre> LOG ACTIVITY LIST----- DB/WORKBENCH --- DB2=D41A----- MEMBER 1 OF 72 Cmd ==>> Scroll ==>> PAGE Cmds: FILT Sels: ? (Menu) B Browse E Edit D Delete P PRNT X Re-execute ----- SEL DATE --- TIME FUNCTION DB2 USERID RE-EXEC MEMBER ----- - 11/07/99 10:05 QKCPROCL DSN2 GBECK2 Q LG000252 - 11/06/99 23:34 QKCPROCL DSN2 GS002 Q LG000251 - 11/06/99 23:31 QKCPROCL DSN2 GS002 Q LG000250 - 11/06/99 18:41 QKCPROCL DSN GS002 Q LG000249 - 11/06/99 18:31 QKCPROCL DSN GS002 Q LG000248 - 11/06/99 18:30 DSN GBECK3 RCODE S LG000247 - 11/06/99 8:28 DSN GBECK3 RCODE S LG000246 - 11/06/99 5:11 QKCPROCL DSN2 GBECK2 Q LG000245 - 11/02/99 18:49 ABEND DSN2 JWHEA 000008 LG000244 - 10/31/99 18:45 ABEND DSN2 JWHEA 0C4000 LG000243 - 10/24/99 17:06 ABEND DSN2 TDD848 000008 LG000242 - 10/24/99 16:42 QKCPROCL DSN2 TDD848 Q LG000241 - 10/23/99 16:17 QKCPROCL DSN2 TDTDB3 Q LG000240 - 10/23/99 15:10 QKCPROCL DSN2 TDTDB3 Q LG000239 - 10/23/99 14:53 QKCPROCL DSN2 TDTDB3 Q LG000238 </pre> </div>

Field definitions for the Log Activity List panel

Each entry in the activity log displays the following information.

Field Name	Description
DATE --- TIME	Date and time function was executed.
FUNCTION	Type of function executed.
DB2	Identifier of the DB2 system on which the function was performed.
USERID	TSO user who executed the function.
RE-EXEC	Code used to re-execute member.
MEMBER	Name of Log PDS member to which the logging information was written.

Methods for manipulating the activity log members

Review the chart for methods to manipulate the activity log members.

Task You Want to Perform	Method	Name
Display the contents of a member	Select	B
Delete a member	Select	D
Edit the contents of a member in an ISPF edit session	Select	E
Print a member	Select	P
Re-execute the statements contained in a member	Select	X

If the RE-EXEC column contains an RCODE value, you can re-execute the action. The value associated with the RCODE indicates the contents of the member to be re-executed:

- RCODE S indicates the member contains JCL
- RCODE R indicates the member contains SQL statements

Option 8: Display the !DB/Tools Allocated Profile Datasets

Overview

This unit contains introductory information on displaying the !DB/Tools allocated profile datasets.

Background about displaying the !DB/Tools allocated profile datasets

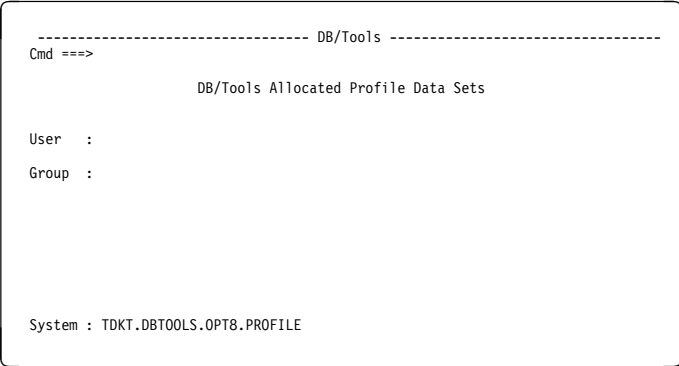
The !DB/Tools profile datasets display panel shows the profile datasets that are currently active. You cannot change the profile datasets on this panel.

Authority required to display the !DB/Tools allocated profile datasets

To use the option for displaying the !DB/Tools profile datasets, you must have at least R/O access to the profile dataset.

Displaying the !DB/Tools Allocated Profile Datasets panel

Follow these steps to display the !DB/Tools Allocated Profile Datasets panel.

Step	Action
1	On the Housekeeping options menu, type 8 in the Option field.
2	Press Enter. Result: The !DB/Tools Profile Datasets are displayed.  <pre> ----- DB/Tools ----- Cmd ==> DB/Tools Allocated Profile Data Sets User : Group : System : TDKT.DBTOOLS.OPT8.PROFILE </pre>

Field definitions for the !DB/Tools Allocated Profile Datasets panel

You cannot change the values for the fields displayed on the panel.

Field Name	Description
User	The !DB/Tools User Profile dataset name (if allocated).
Group	The !DB/Tools Group Profile dataset names (if allocated).
System	The !DB/Tools System Profile dataset name.

Option 9: Set the !DB/Tools Profile Datasets

Overview

This unit contains introductory information on setting the !DB/WORKBENCH profile datasets.

Background about setting the !DB/Tools profile datasets

The !DB/Tools Profile Information panel displays the profile datasets that are currently active, allowing you to set the profile datasets and the order in which they will be used.

Authority required for the !DB/Tools profile datasets

To use this option for the !DB/Tools profile datasets, you must have these authorizations:

- To display information only, you must have at least R/O access to the profile dataset.
- To change values, you must have R/W access to the profile dataset.

Setting the !DB/Tools Profile Datasets

Follow these steps to set the !DB/Tools profile datasets.

Step	Action
1	On the Housekeeping options menu, type 9 in the Option field.
2	<p>Press Enter.</p> <p>Result: The !DB/Tools Profile Information panel is displayed.</p> <div data-bbox="573 573 1250 989" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> ----- DB/Tools ----- Cmd ==> DB/Tools Profile Information Save Values? ==> Y (Y Permanent N Temporary Display this panel on product entry? ==> N (Y Yes N No) Use? Type Profile Data Set Name Status ----- N User : _____ N Group: _____ N _____ N _____ N _____ N _____ N _____ N _____ N _____ Y System: TDKT.DBTOOLS.OPT8.PROFILE OK </pre> </div>
3	You can change the values for the fields displayed on the !DB/Tools Profile Information panel.

Field definitions for the !DB/Tools Profile Information panel

Field Name	Description
Save Values?	Specify Y if you want any changes made on this panel to be saved permanently in your current ISPF profile. Specify N if you want any changes made on this panel to only be used for the current !DB/Tools session.
Display this panel on product entry?	Specify Y if you want this panel displayed every time you enter a !DB/Tools product. Specify N if you do not want this panel displayed.
Use?	Specify Y if you want to use the specified profile dataset when accessing !DB/Tools. Specify N if you do not wish to use this profile data set.
Type	The profile dataset type (User, Group, or System).
Profile Dataset Name	The name of the !DB/Tools profile dataset. The system profile data set is defined during the !DB/Tools installation and cannot be changed from this panel.
Status	The current Status of the profile dataset. (For a description of the possible statuses of the PROFILE datasets see “Status definitions for the !DB/Tools profile datasets” on page 194.)

Status definitions for the !DB/Tools profile datasets

Status	Description
OK	Profile dataset exists and has the current dataset attributes.
Migrated	Profile dataset is migrated.
INVALID DSORG	Profile dataset exists but is not a partitioned dataset.
INVALID RECFM	Profile dataset exists but is not a fixed block dataset.
Dataset NOT FOUND	Profile dataset does not exist.

Precedence of the profile datasets

The product uses the profile datasets in the order they are displayed on the panel. If the member the product needs is not in the first profile dataset displayed on the panel, the product checks for the member in the next profile dataset displayed on the panel.

For example, if you specify the name of a user profile dataset, the product checks for the member in the user profile dataset first. If you do not specify the name of a user profile dataset, the product checks for the member in the group profile dataset first.

Chapter 7. Managing Statistics in the STATDBAS Table Space

Introduction

This chapter provides general background about the statistics that are available from the tables in the STATDBAS table space that you can use to run queries. It includes instructions for managing these tables and covers the contents of each table.

Chapter contents

Overview of the Chapter	196
General Background about the STATDBAS Table Space	198
Background about Growth in Size of Tables	201
Controlling the Size of Tables	203
Contents of the Tables in STATDBAS Table Space	204

Overview of the Chapter

Background about the extract and the STATDBAS table space

!DB/WORKBENCH allows you to run queries against statistical data that !DB/WORKBENCH inserts in the tables in the STATDBAS table space. When you run an extract, !DB/WORKBENCH automatically builds DB2 INSERT statements for each table space, table, and index with statistics. (The system created these statistics when you used the RUNSTATS utility.)

This chapter covers issues related to the STATDBAS table space and to controlling the inserting of statistical data in the tables in the table space. For information about the extract itself, see the chapter “!DB/WORKBENCH Housekeeping Options” on page 139.

Organization of information in this chapter

This chapter covers background about the connections between the simultaneous processes of running an extract and of providing selected data for the tables in the STATDBAS table space. The chapter also covers the purpose and the operation of the tables in the STATDBAS table space and contains information about how to control the contents of the tables in the STATDBAS table space. The chapter concludes with detailed information about the types of data included in the tables in the STATDBAS table space.

Organization of the chapter and your needs

Review the chart to determine the units you need to use.

	General Background	Growth	Controlling	Contents of Tables
	p.198	p.201	p.203	p.204
Review background about the tables and their operation in the STATDBAS table space	√			
Review background about possible growth in the size of the tables in the STATDBAS table space	√	√		
Reduce the number of rows of data in the tables (You want to continue to be able to run queries against statistics in the tables in the STATDBAS table space.)	√	√	√	
Perform actions so that !DB/WORKBENCH no longer collects statistical data (You no longer want to be able to run queries against statistics in the tables in the STATDBAS table space.)	√	√	√	
Review contents of the tables in the STATDBAS table space				√

General Background about the STATDBAS Table Space

Overview

This unit contains general background about the tables in the STATDBAS table space and about the relationship between the STATDBAS table space and running of extracts. It includes a summary showing the tables in the STATDBAS table spaces that are required to run !DB/WORKBENCH and the tables that are required to use the statistical data.

Background about the STATDBAS table space

The STATDBAS table space is a DB2 table space in database KTCDBD (a database that you can rename). The table space contains tables for several products; however, these tables are specific to !DB/WORKBENCH:

- SYSTABLESPACE
- SYSTABLES
- SYSINDEXES

Relationship between the STATDBAS table space and extracts

When you run an extract, !DB/WORKBENCH stores the extract in the system PDS. !DB/WORKBENCH also performs this set of actions with the statistical data and the tables in the STATDBAS table space:

- Builds DB2 INSERT statements for each table space, table, and index that has statistics (which were created when you used the RUNSTATS utility)
- Applies a time stamp to the data (This information is stored in each table under the column heading *STATTIME*.)
- Executes the INSERT statements and inserts the rows of data in the appropriate table (For example, it inserts rows of data about table spaces in the SYSTABLESPACE table.)

How to accumulate statistics in the STATDBAS table space

If you want to accumulate statistics, you must specify at installation time that you want to gather statistics. Also, at installation time, if you want statistics, you must specify the number of days of statistics to keep. Statistics will be accumulated in the STATDBAS table space. The STATDBAS table space is not necessary to run !DB/WORKBENCH, but it is necessary if you want to run queries against the statistical data in the tables in the STATDBAS table space.

Summary of the requirements for tables in the STATDBAS table space

The chart provides a summary showing the tables required by !DB/WORKBENCH and the tables required if you want to query statistics.

Table Name	Required by !DB/WORKBENCH	Required if You Want Statistics
SYSINDEXES	No	Yes
SYSTABLES	No	Yes
SYSTABLESPACE	No	Yes

How to stop accumulating statistics in the STATDBAS table space

Each table in the STATDBAS table space has a prefix of *<creatorid>* that was established at install time. The name of each of the tables within the STATDBAS table space is:

- *<creatorid>*.SYSTABLESPACE
- *<creatorid>*.SYSTABLES
- *<creatorid>*.SYSINDEXES

If you want to stop accumulating statistics in the STATDBAS, drop any one or all of the tables in the STATDBAS table space. If you want to drop only one, the *<creatorid>*.SYSTABLESPACE would be best choice since it is the first table to be updated.

How to clean up the STATDBAS table space

When you run the install, the following STEP50 will be included in the extract JCL if you specify STATS=Y at install time. Statistics will be kept for the number of days you specify at install time.

If you have been accumulating statistics in the STATDBAS table space, and you want to clean up the tables within it, you can rerun the portion of the install that cleans up the STATDBAS table space or you can write your own program to accomplish this task. Add the JCL in the following example, named STEP50, to your program to purge data from the STATDBAS table space. When you add the JCL in STEP50 to your routine, fill in the table owner in the *<creatorid>* field for each table in the example. Also, modify the number of days you want to keep the statistics in each of the tables (60 days is provided in the example for each table).

```
//STEP50    EXEC= PGM=IKJEF01,REGION=2M,DYNAMNBR=100
//STEPLIB  DD DISP=SHR,DSN=PP.DB2.V3R1.SDSNLOAD
//          DD DISP=SHR,DSN=PP.DB2.V310.RUNLIB.LOAD
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *
            DSN SYSTEM(DB31)
            RUN PROGRAM(DSNTEP2) PLAN(DSNTEP31)
            END
//SYSIN DD *
DELETE FROM <creatorid>.SYSTABLES
WHERE DAYS(CURRENT TIMESTAMP) - DAYS(STATTIME) > 60;
DELETE FROM <creatorid>.SYSINDEXES
WHERE DAYS(CURRENT TIMESTAMP) - DAYS(STATTIME) > 60;
DELETE FROM <creatorid>.SYSTABLESPACE
WHERE DAYS(CURRENT TIMESTAMP) - DAYS(STATTIME) > 60;
```

How to reinstate the STATDBAS table space

If you stopped accumulating statistics or if during installation you elected not to gather statistics in the STATDBAS table space, you may reinstate accumulating statistics by going through the build portion of the installation process that creates the tables in the STATDBAS table space and specifying you want statistics (STATS=Y).

Background about Growth in Size of Tables

Overview

This unit provides necessary background and an example of how the size of tables can grow each time you run an extract if you either have not made changes to the system during installation or if you do not perform the actions in the unit “Controlling the Size of Tables” on page 203. The unit “General Background about the STATDBAS Table Space” on page 198 is a prerequisite for this unit.

Background about Lazarus Retention and extracts

You can use the Lazarus Retention field available from the !DB/WORKBENCH Housekeeping menu to specify the number of extracts that !DB/WORKBENCH stores in the system PDS. The Lazarus Retention field, however, has no effect on the retention of the inserted rows of data in SYSTABLESPACE, SYSTABLES, or SYSINDEXES tables.

Characteristics of the example

These characteristics exist for the example shown in this unit:

- You have not taken action during installation to prevent the system from saving statistical data to the STATDBAS table space.
- You have specified the Lazarus Retention value to be 4 using option **3** on the !DB/WORKBENCH Profile Data Administration menu. (This menu is accessed by typing **1** on the Housekeeping menu.) (The system will, therefore, retain only the four most recent extracts.)
- You are running your fifth extract.

Note: If, during the install process, you specify that you do not want to collect statistics during an extract, !DB/WORKBENCH returns an informational code of 4 from the extract process.

Consequences of the example

Review the chart to see the response of the system with the extract and with the tables SYSTABLESPACE, SYSTABLES, and SYSINDEXES. (The Example column in the chart is based on the characteristics described in this unit.)

Location	Action by the System	Example
In the system PDS	The system stores the new extract in the system PDS and manages the number of extracts according to the value you set for the Lazarus Retention field.	The system saves the new extract to the system PDS and also deletes the first extract. At the end of the operation, the system PDS contains four (4) extracts.
In the SYSTABLESPACE, SYSTABLES, and SYSINDEXES tables in the STATDBAS table space	The system time stamps the data that has statistics for table spaces, tables, indexes, and columns and inserts the rows in the appropriate tables.	The system inserts the rows of data in the appropriate tables. At the end of the operation, the SYSTABLESPACE, SYSTABLES, and SYSINDEXES tables each contain a total of five (5) sets of statistical data.

Controlling the Size of Tables

Overview

This unit provides instructions for controlling the size of the tables in the STATDBAS table space. The instructions vary depending upon whether you want to use statistics or whether you no longer want to collect statistics.

Reminder about the STATTIME column and identifying outdated data

You can use the time stamp value in the STATTIME column to determine which rows in the table are outdated. If you need a reminder about the time stamp displayed in this column, see “General Background about the STATDBAS Table Space” on page 198.

Methods for controlling the size of the tables

Review the chart for the methods appropriate for you.

IF . . .	THEN . . .
You want to continue to use the statistics, but you need to reduce the size of the tables and to remove outdated data	Delete any outdated data contained in the SYSTABLESPACE, SYSTABLES, and SYSINDEXES tables by writing an SQL query.
You no longer want to use statistics that are available with the tables in the STATDBAS table space	Delete the contents of the SYSTABLESPACE, SYSTABLES, and SYSINDEXES tables. Allocate a PRIQTY value of 1 and a SECQTY value of 0 for the tables.

Contents of the Tables in STATDBAS Table Space

Overview

This unit contains a series of charts showing the columns for each of the four tables in the STATDBAS table space that contain statistics related to !DB/WORKBENCH. (The charts show the name of the database to be KTCDBD; however, you can rename the database.)

Reminder about the STATTIME column

To distinguish between data inserted during different extracts, see the STATTIME column. If you need a reminder about the time stamp displayed in this column, see “General Background about the STATDBAS Table Space” on page 198.

Table KTCDBD.SYSTABLESPACE in KTCDBD.STATDBAS

The chart shows the type, length, and description for each column in the table KTCDBD.SYSTABLESPACE.

COLUMN NAME	TYPE	LEN	DESCRIPTION
DBNAME	CHAR	8	Database Name
TSNAME	CHAR	8	Table Space Name
TSCREATOR	CHAR	8	Table Space Creator
PARTITION	SMALLINT	2	Partition Number
STATTIME	TIMESTAMP	10	Time Statistics Added
NACTIVE	INTEGER	4	Number Pages In Use
SPACE	INTEGER	4	Total K in Space
CARD	INTEGER	4	Number Rows In Table Space
FARINDREF	INTEGER	4	Number Far from Home
NEARINDREF	INTEGER	4	Number Moved Nearby
PERCACTIVE	SMALLINT	2	Percent Pages Active
PERCDROP	SMALLINT	2	Percent Pages Dropped

Table KTCDBD.SYSTABLES in KTCDBD.STATDBAS

The chart shows the type, length, and description for each column in the table KTCDBD.SYSTABLES.

COLUMN NAME	TYPE	LEN	DESCRIPTION
DBNAME	CHAR	8	Database Name
TSNAME	CHAR	8	Table Space Name
TSCREATOR	CHAR	8	Table Space Creator
PARTITION	SMALLINT	2	Partition Number
STATTIME	TIMESTMP	10	Time Statistics Added
CARD	INTEGER	4	Number Rows in Table
NPAGES	INTEGER	4	Number Pages with Rows
PCTPAGES	SMALLINT	2	Percent Pages with Rows
RECLENGTH	SMALLINT	2	Computed maximum length of any record in the table
TBCREATOR	CHAR	8	Table Creator ID
TBNAME	VARCHAR	18	Table Name

Table KTCDBD.SYSINDEXES in KTCDBD.STATDBAS

The chart shows the type, length, and description for each column in the table KTCDBD.SYSINDEXES.

COLUMN NAME	TYPE	LEN	DESCRIPTION
DBNAME	CHAR	8	Database Name
TSNAME	CHAR	8	Table Space Name
TSCREATOR	CHAR	8	Table Space Creator
PARTITION	SMALLINT	2	Partition Number
STATTIME	TIMESTAMP	10	Time Statistics Added
INDEXSPACE	CHAR	8	Index Space Name
CLUSTERING	CHAR	1	Clustering Index?
CLUSTERED	CHAR	1	Clustered Index?
FIRSTKEYCARD	INTEGER	4	Number Values of Key Field
FULLKEYCARD	INTEGER	4	Number Values of Key
NLEAF	INTEGER	4	Number Leaf Pages
NLEVELS	SMALLINT	2	Number Index Levels
CLUSTERRATIO	SMALLINT	2	Percent Index Clustered
SPACE	INTEGER	4	Total K in Space
CARD	INTEGER	4	Number Rows Indexed
FAROFFPOS	INTEGER	4	Number Far from Home
LEAFDIST	INTEGER	4	100 x Average Distance
NEAROFFPOS	INTEGER	4	Number Move Nearby
IXCREATOR	CHAR	8	Index Creator ID
IXNAME	VARCHAR	18	Index Name
TBCREATOR	CHAR	8	Table Creator ID
TBNAME	VARCHAR	18	Table Name

Creating Objects

Chapter 8. Introduction to Creating Objects

Introduction

This chapter introduces the section on using the Create Menu to generate statements to create objects and, optionally, to define the VSAM datasets for table spaces or index spaces. The chapter is a prerequisite for the other chapters in this section.

Chapter contents

Resources for Learning about Topics in This Section	210
Overview of Chapter	212
Overview of Methods to Create Objects	214
How to Access and Select the Option You Want	216
How to Determine Order of Creation of Objects	218
Methods Available If You Make a Mistake	219
Introduction to Methods for Providing Values	220
Methods Available to Estimate Requirements for Space	222
Overview of Operation of Defaults	223
Background about Memory and Options	226
Memory and Effects on Selecting from Lists	227
Memory and Effects on Ending Sessions	228
Memory and Effects on Beginning a Session	230
Processing the Statements	232

Resources for Learning about Topics in This Section

Overview

This unit lists resources available for learning about options for creating objects from the Create Menu. It includes information about the resources in this guide.

Resources in this guide

This guide is intended to provide prerequisites and instructions for the variety of methods !DB/WORKBENCH provides for creating objects. The chart lists sources available within this section.

Information You Want	Resources Available
Required introduction to use any of the options on the Create Menu (This chapter is a prerequisite for the other chapters in this section.)	This chapter
Prerequisites and instructions for different methods for creating a database	“Creating a Database” on page 235
Prerequisites and instructions for different methods for creating a table space (This chapter also includes instructions for DB2-created and managed VSAM datasets for a non-partitioned table space.)	“Creating a Table Space” on page 247
Prerequisites and instructions for creating a table	“Creating a Table” on page 265
Prerequisites and instructions for different methods for creating an index and index space (This chapter also includes the instructions for DB2-created and managed VSAM datasets for a non-partitioned index space.)	“Creating an Index and Index Space” on page 283
Instructions for completing partitioned table space and index spaces and for completing other options for VSAM datasets	“Completing Requirements for VSAM Datasets” on page 315

Resources in this guide on values needed

This guide includes information about the values that you must provide when creating objects. For information see “Values Used When Creating or Altering Objects” in *!DB/WORKBENCH Reference Manual*.

Resources in online Help

If you are providing values in fields on a panel and want immediate information about possible values, type **HELP** on the command line and press Enter, or press the appropriate function key.

Overview of Chapter

Background about !DB/WORKBENCH and this section

Creation of objects required in database administration is time-consuming and error-prone. !DB/WORKBENCH provides an extensive facility for creating databases, table spaces, tables, indexes, or storage groups. When you create a table, you can also create synonyms. When you create table spaces or index spaces, !DB/WORKBENCH also supports the defining of implicit or explicit VSAM datasets.

This section covers the prerequisites and instructions you need to be able to create objects. This chapter provides the prerequisite instructions for the entire section.

Organization of information in the chapter

The chapter introduces the methods you can use to create objects with !DB/WORKBENCH. It also contains information essential to the other chapters in this section. It also contains instructions for accessing the options on the Create Menu.

Organization of the chapters in this section

The chapters for creating objects follow the same organization:

- An introduction to how you can access the first panel for the option (and instructions if you want to exit)
- Prerequisites and authority required for creating the objects
- Instructions for beginning the process and an overview of the methods possible for creating the objects and the units to use for the method. (With tables, there is only one method for creation.)
- Instructions for each of the methods

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Overview	Access and Select	Order of Creation	Recovery	Values	Estimator	Defaults	Background	Selection	Ending	Beginning	Processing
	p.214	p.216	p.218	p.219	p.220	p.222	p.223	p.226	p.227	p.228	p.230	p.232
Determine which method for generating CREATE statements matches your needs	√											
Review options on the Create Menu and determine how to access options		√										
Determine an order for creating the objects you want			√									
Review methods for recovery from error				√								
Review how you can reduce keying when providing values in general and when using lists					√							
Review an introduction to using the DASD space estimator when creating a table space or index space						√						
Use default values available with !DB/WORKBENCH							√					
Review how memory affects using options on the Create Menu								√	√	√	√	
Review methods for processing and executing statements												√

Overview of Methods to Create Objects

Overview

This unit provides a comparison of the methods for creating objects with !DB/WORKBENCH. Review the unit if you want to determine which of the methods is best for your needs or if you have previously used one method and want to be aware of differences in the methods.

Background about sources

If you want more information about the OUT command, the OUTC command, or the Out DDL (O) select, see the chapter “Methods for Generating Statements for Objects” in *!DB/WORKBENCH Basic User's Guide*.

Methods for creating objects

The chart illustrates the CREATE statements that result from the methods.

CREATE Statements Generated	Method for Creating Object
For all objects listed and for all the dependent objects	OUT command
For all objects listed and for all dependent objects you specify on the OutDDL Options panel	OUTC command
For the selected object and for its dependents only	Out DDL (O) select
For a single object (The CREATE statements do not include those for dependent objects.)	Option on the Create Menu

Methods for storing or saving generated statements

The chart shows the differences in storing or saving statements.

Method for Creating Objects	Output of Statement
OUT and OUTC commands and the Out DDL (O) select	In a member of the !DB/WORKBENCH user PDS
Options from the Create Menu	During the session, in memory
	After a session, in a member of the !DB/WORKBENCH user PDS

Methods for keying of values

The chart shows how the methods compare on ways to reduce keying of values required for CREATE statements. (For example, !DB/WORKBENCH provides the syntax for statements, but you must specify the values to use in the statements, such as the name of the storage group and the value for the ROSHARE parameter, if used.)

Method for Creating Objects	Sources for Values
OUT and OUTC commands and the Out DDL (O) select	Existing objects (selected from an object list panel) that serve as models for the values required
Options from the Create Menu	Existing objects (selected from a list) that serve as models for the values required
	Default values that you provide using the default option on the Create Menu
	Entry of values on the panel

Method if you want to generate statements for referential integrity

If you want to generate statements for referential integrity (such as ALTER statements used to establish primary and foreign keys), you use the OUTC command. With the OUTC command, you select an object with referential relationships and use the OUTC command to display the OutDDL Options panel. On that panel, you can specify that you want the option Refs. The system generates only the ALTER statements necessary to establish referential integrity.

With RFO command, you can also generate ALTER statements for existing objects. If you need a reminder about the RFO command, see the chapter “Methods for Generating Statements for Objects” in *!DB/WORKBENCH Basic User's Guide*.

Method if you want to define a VSAM dataset at creation

If you want to define a VSAM dataset at the same time as you create the table space or index space, you use the options available from the Create Menu.

How to Access and Select the Option You Want

Overview

This unit contains an overview of options you can select and instructions for how you can access the Create Menu or simultaneously access the Create Menu and one of its options.

Background about methods for defining VSAM datasets

When you choose to create either a table space or index, you also must define the VSAM datasets. !DB/WORKBENCH supports three methods:

- A DB2-created and managed dataset
 - An existing VSAM dataset that you will manage
 - A VSAM dataset that you will manage and that you will define during this session
-

Options for creating objects on the Create Menu

The chart lists the objects you can create and the option to select.

Object You Want to Create	Option to Use
Databases	D
Table spaces (and define the VSAM datasets)	S
Tables (and synonyms on a table)	T
Index (and index space and define the VSAM datasets)	I
Storage group	G

Procedure to access the Create Menu

Follow these instructions to access the Create Menu.

Step	Action
1	On the !DB/WORKBENCH Primary Menu, type CR in the Option field.
2	Press Enter. Result: The system displays the Create Menu.

Procedure to access the Create Menu and an option

If you know the letter of the option on the menu, you can select both the Create Menu and the option. Follow these steps.

Step	Action
1	On the !DB/WORKBENCH Primary Menu, type CR , a period (.), followed by the character for the option. Example: Type CR.D to access the Create Menu and the Create Database option.
2	Press Enter.

How to Determine Order of Creation of Objects

Overview

This unit provides background about how the system operates with dependent objects and suggestions for the order in which you create objects.

Background about dependent objects

If you type the name of an owning object that is not defined in the DB2 catalog, the system automatically displays the panel to create the object. For example, if you provide the name of a new database on the Create Table Spaces panel, the system automatically displays the Create Databases panel. The chart shows your options.

Conditions	Action to Perform
You accidentally typed a new name and you do not want to create a new object	Blank out the new name on the panel and use the END command. Result: The system redisplay the panel where you mistyped the name.
You want to create a new object	Complete all panels required for the owning object. Result: The system redisplay the panel where you typed the name of the new object.

Order of creation of objects

The chart shows the objects that are required to create another object. If you are creating new objects with dependents, you can create the objects in the order shown on the Create Menu.

Object You Want to Create	Prerequisite Object
Database	Storage group or equivalent VOLSERS (or you can select the default SYSDEFLT from the list)
Table space	Database
Table	Database and table space
Index	Database, table space, and table

Methods Available If You Make a Mistake

Overview

This unit reviews methods available in !DB/WORKBENCH to recover from a mistake.

Redisplaying the previous panel

!DB/WORKBENCH provides the REV command on some panels so that you can redisplay the previous panel in a series of panels. (The panel includes a statement if you can use the command. In most cases, you cannot redisplay the previous panel when specifying values for partitions.)

Task You Want to Perform	Action
Void all values on the current panel, and redisplay the previous panel	Type REV on the command line and press Enter.

Stopping the current option

If you are using an option to create an object and decide that you no longer want to create the object, you can use the command shown in the chart.

Task You Want to Perform	Action
Void all values on all panels used with the option (The system voids only the panels for the object you are currently creating; it does not void any objects you have completed creating.)	<p>On the panel used with the Create option, type END on the command line and press Enter, or press the appropriate function key.</p> <p>Result: The system displays a panel with instructions for stopping the option or returning to it.</p> <p>On the panel with instructions, type END on the command line and press Enter, or press the appropriate function key.</p> <p>Result: The system returns to prior menu or panel.</p>

Introduction to Methods for Providing Values

Overview

This unit introduces methods for providing values needed in statements and reviews the methods available in !DB/WORKBENCH for selection from lists and for filtering lists.

Methods available that you can use to provide values

When you use the Create Menu, !DB/WORKBENCH generates the CREATE statements and provides the correct syntax; however, you must specify the values to be used in these statements. (For example, if you want to create a database, you must specify the name of the buffer pool and the storage group and the value for the ROSHARE parameter.) !DB/WORKBENCH supports three methods for you to use to provide these values:

- Select an object as a model from a list of existing objects (The system copies the values from the object into the fields on the panels and, subsequently, into the generated statements.)
 - Use default values that you specify using the default option (O) on the Create Menu or that !DB/WORKBENCH provides. (The defaults are particularly helpful in reducing keying for table spaces and index spaces.)
 - Type the values directly (You can also type over values; for example, you can select a database as a model but type over the value displayed for buffer pool.)
-

Methods for quick selections from lists

!DB/WORKBENCH supports the LOCATE command and the S select to specify the object you want to use as a model.

Task You Want to Perform	Action
Move the cursor on the panel to the name of a specific object	Type L followed by a space and the name of the object on the command line and press Enter.
Select an object to use as a model	Type S in the Select field for the object and press Enter.

Filtering and selections from lists

You can use filtering to specify objects that you can select as a model for the object you are creating or as the name for an object that you must specify while creating another object. (For example, you must specify the owning database and table space when you create a table.)

When you use the FILT command from any panel, the system displays the FILT selection list for the Table Spaces panel. You can use the filters on the selection list, or you can use one of the listed filters directly on the command line. You can also use the NOFILT command to turn off filtering.

Other methods to filter data

!DB/WORKBENCH also provides additional methods to filter data on some panels.

Example	Information You Want to Display	Method to Use
The table spaces within a database	The objects associated with another object	* (asterisk) in the field for the object
The storage groups available for selection	The list of objects available for the specified object	? (question mark) in the field for the object

Reminder about filters

If you need a reminder about filtering, see the chapter “Filtering Object Lists” on page 195.

Methods Available to Estimate Requirements for Space

Overview

This unit provides a brief introduction and an overview of the procedures to use with the Space Estimator that you can use when defining VSAM datasets when creating a table space or index space.

How to use the Space Estimator when creating objects

You can use the Space Estimator on a panel with the field for entry of Priqty. (If you can invoke the panel during creation of a table space or index space, the panel includes a phrase "?" *to Est.*) For more information about this estimator, see the chapter "Using the DASD Estimators" on page 111.

Step	Action
1	On the panel listing values, type ? in the Priqty field.
2	Press Enter. Result: The system displays the table space Estimator panel.
3	Type values you want to use in the estimator.
4	Press Enter. Result: The system displays the results of the calculations by the estimator.
5	Press END. Result: The system redisplay the original panel and lists the output from the estimator in the Priqty field.

Overview of Operation of Defaults

Overview

This unit contains background on using defaults, the conditions for overriding defaults, the values you can set on the default panel, and additional information about these defaults.

Background about defaults and creating objects

When you use the default option (option O) on the Create Menu, the system saves the defaults in the ISPF profile. If you do not specify a default value on the panel, you can select the default from the lists that !DB/WORKBENCH displays. For example, if you are creating a storage group, you can select SYSDEFLT from the list that !DB/WORKBENCH displays. The values that !DB/WORKBENCH provides are not necessarily the IBM standard DB2 defaults.

Default values and other sources for value

You can change defaults directly by typing a different value on the default panel. You also can indirectly and temporarily change the values that the system displays in fields during a session by performing either of these actions:

- By typing a different value over an existing default value
 - By selecting as the model for the object an object that has values different from the defaults
-

Restoring default values

When you overwrite values by either typing over the values or selecting a model, you cannot restore the default values. If you wanted to restore these values, you would need to exit the Create Menu and return.

Defaults available

This chart lists the types of defaults available when you select the option for default on the Create Menu.

Type of Default	Field Name
To control the timing of display of panels when you provide the name of a new object when you are creating another object	New Object Waiting Return Wait
To identify the names of objects that you will use when creating other objects	Database Storage Group
To specify values used for several objects	Buffer Pool Creator ID
To specify the values for all types of table spaces and index spaces	CLOSE FREEPAGE Gbpcache LOCK PCTFREE
To specify values specific to table spaces	COMPRESS
To specify values specific to indexes	DEFER UNIQUE
To specify the values for VSAM datasets that DB2 manages and creates or that you plan to manage and to create during this session	ERASE PRIQTY SECQTY
To specify the values for existing VSAM datasets that you manage (also used on the Create Storage Group panel)	VSAM Catalog Password
To cache changed pages into the group buffer pool as they are read from DASD.	Gbpcache

Additional information about defaults for timing of displays

You cannot create a dependent object without specifying the owning objects. For example, you cannot specify a table without specifying its owning table space. If you provide the name of a new table space when you are creating a table, the system automatically displays the panel so you can create the table space.

With the option, you can control two types of timing of displays of panels:

- The time from your entry of the new name and to the display of the create panel
- The time from your completion of the create panel to the redisplay of its original panel

Additional information about the DEFER value

If you want to use the DEFER option with indexes, perform either of these actions to create the index.

- Use the CR option on the !DB/WORKBENCH Primary Menu
- Use the Out DDL (O) select, the OUT command, or the OUTC command, and then edit the contents of the PDS member to add the DEFER clause.

Information about values for defaults

If you want information about values appropriate for the defaults that you can specify on the panel, see the chapter “Values Used When Creating or Altering Objects” in *!DB/WORKBENCH Reference Manual*.

Background about Memory and Options

Overview

This unit provides background information about use of memory and the options available for creating objects. It is a prerequisite for the three units that follow.

How statements are stored when you use options

During the session when you create objects using any of the options available on the Create Menu, the system stores the CREATE statements for the objects (including the values you specify) in memory. It does not save these statements to a member of the !DB/WORKBENCH user PDS until the end of the session.

The chart illustrates the beginning and ending of sessions and the time when the system uses memory to store statements.

Your Action	System Response
Select an option from the Create Menu and begin to create an object	Stores all values in memory
Return to the Create Menu and type END on the command line and press ENTER, or press the appropriate function key	Ends the session, saves the statement to the <i>useridC</i> member of the !DB/WORKBENCH, and displays a version of the Process Member menu for you to use to specify what you want to do with the statement

Benefits of storage of statements in memory during a session

!DB/WORKBENCH stores the create statements in memory (rather than saving them to a member) so that you have the choice of creating one or more dependents for a new object as well as the new object. For example, when you create a new object (such as a table space), you can also create a dependent (such as a table).

Memory and Effects on Selecting from Lists

Overview

This unit discusses how the use of memory with the Create options affect interpreting lists of selections and how it affects creation of dependent objects. The unit “Background about Memory and Options” on page 226 is a prerequisite for this unit.

Implication of memory on interpreting lists of selections

When you create a new object, the system stores the CREATE statements in memory. When the system subsequently displays lists of that type of object, it displays the names of these objects:

- The new name of the object from memory (displayed in reverse video)
- The name of existing objects from the extract

For example, if you select the Table Spaces option on the Create Menu and complete all panels needed to create a table space that you name TS1234, the system displays a panel so that you can choose to create a dependent table without returning to the Create Menu. If you choose to create the table, the system lists the names of the tables. When the system displays the list of possible owning table spaces for the new table, the first name on the list is table space TS1234 (displayed in reverse video).

Implications for creating dependent objects

If you want to create statements for dependent objects using the options on the Create Menu, do not end the session until you create all of the dependent objects that you want. Once you end the session, the system saves the statements to a member of the !DB/WORKBENCH user PDS and, therefore, does not display the objects on selection lists.

For example, if you create table space TS1234 and a table TB9999 and if you end the session, the system saves the CREATE statements from memory to the !DB/WORKBENCH user PDS member (*useridC*). If you select the Create option from the !DB/WORKBENCH Primary Menu and then the Table option from the Create Menu, the system does not display table space TS1234 on the lists for selection of owning table spaces.

Memory and Effects on Ending Sessions

Overview

This unit discusses implications of how memory affects the ending of a session. The unit “Background about Memory and Options” on page 226 is a prerequisite for this unit.

Memory and ending a session

When you end the session of creating objects, the system saves the CREATE statements (including values) from memory to the *useridC* member in the !DB/WORKBENCH user PDS. The system also displays a Process Member menu.

```

PROCESS MEMBER ----- DB/WORKBENCH -- DB2=D23A -----
Cmd ==>                                     99/10/28 06:26

The CREATE SQL statements have been saved in member TDTD86AC.

Option ==>

      1 UPDATE NOW      Invoke DB2 to execute the statements.
      2 EDIT            Edit saved SQL statements.
      3 DELETE          Delete saved SQL statements member.
1  4 PRINT            Print the SQL statements member.
      5 UPDATE ACCUM    Add these stmts to the Accum member.
      6 START ACCUM     Start a new Accum Member.
      7 LOG             Copy the member to the Log PDS.

2  8 DO LATER        Changes are saved, exit now.

```

1 Options available from many Process Member menus

2 Option distinctive to ending a session from the Create Menu

Additional information about commonly available options

If you need a reminder about commonly available options on the Process Member menu, see the chapter “Introduction to Operations on Data” in *!DB/WORKBENCH Basic User's Guide*.

Additional information about the DO LATER option

If you select the DO LATER option or use the END command, the system performs these actions:

- Leaves the member *useridC* in the !DB/WORKBENCH user PDS (The system does not overwrite the member.)
- Displays the !DB/WORKBENCH Primary Menu

Implications for beginning another session

If you select the DO LATER option or use the END command when you end a session, the system displays a Process Member menu when you next select the Create option from the !DB/WORKBENCH Primary Menu. For information, see “Memory and Effects on Beginning a Session” on page 230.

Memory and Effects on Beginning a Session

Overview

This unit discusses implications of how memory affects beginning a session. The unit “Background about Memory and Options” on page 226 is a prerequisite for this unit.

Background about memory and beginning the session

When you select an option on the Create Menu and begin to complete panels, the system stores the CREATE statements in memory. The statements remain in memory until you return to the Create Menu and type END on the command line and press Enter, or press the appropriate function key.

Background for beginning a session

The way that you ended your previous session determines the menu that the system displays when you begin a session using the options available on the Create Menu. The chart shows the difference.

Method You Used to End the Previous Session	Menu Displayed
Executed or deleted the <i>useridC</i> member or saved the contents to the Accum member	Create Menu
Selected the DO LATER option or used the END command	A Process Member menu

If the system displays the Create Menu

If the system displays the Create Menu, you can select an option to create an object. (If you want more information about the options, see “How to Access and Select the Option You Want” on page 216.)

If the system displays a version of the Process Member menu

If the system displays a Process Member menu, you must choose from one of the options on the menu.

Interpreting the Process Member menu

If you selected the DO LATER option or used the END command when you ended the session, the system displays a Process Member menu when you begin another session.

```

PROCESS MEMBER ----- DB/WORKBENCH -- DB2=D23A -----
Cmd ==>                                     99/10/28 06:31

The CREATE SQL statements have been saved in member TDTD86AC.

Option ==>

  1  UPDATE NOW      Invoke DB2 to execute the statements.
  2  EDIT            Edit saved SQL statements.
  3  DELETE          Delete saved SQL statements member.
  1 4  PRINT         Print the SQL statements member.
  5  UPDATE ACCUM   Add these stmts to the Accum member.
  6  START ACCUM    Start a new Accum Member.
  7  LOG            Copy the member to the Log PDS.

  2 8  CONTINUE      Prior changes are lost.

```

- 1** Options available from many Process Member menus
- 2** Option distinctive to beginning a session from the Create Menu when you previously entered a session by selecting the DO LATER option

Additional information about commonly available options

If you need a reminder about commonly available options on the Process Member menu, see the chapter “Introduction to Operations on Data” in *!DB/WORKBENCH Basic User's Guide*.

Additional information about the CONTINUE option

If you select the CONTINUE option, the system performs these actions:

- Leaves the member in the !DB/WORKBENCH user PDS (when you end the session the system overwrites the member.)
- Displays the Create Menu

Processing the Statements

Overview

This unit reviews options for processing statements and lists ways to execute these statements from a stored member of the !DB/WORKBENCH user PDS or the statements saved to the Accum member.

Options for processing statements using the Process Member menu

The system displays the Process Member menu when you type END on the command line and press Enter, or press the appropriate function key. The Process Member menu provides options to perform such tasks as execute, edit, delete, or accumulate the statements.

Options for executing statements stored in a member

To execute statements that you have stored in a member of the !DB/WORKBENCH user PDS, follow instructions in the unit “Executing Statements Stored in a PDS Member” on page 38.

Options to execute accumulated statements

You can execute the statements in the member by selecting the UPDATE NOW option, but you also have options for executing the statements in the Accum member.

Condition	Action to Perform
You are on the Process Member menu and you have accumulated statements to the Accum member (The system displays a message indicating the statements are saved to member <i>useridO</i> .)	Type the member of the option UPDATE NOW and press Enter. Result: The system executes the statements in the Accum member.
You are on any object list panel that supports the OUTX command (such as the database panel)	On the object list panel, type OUTX on the command line and press Enter. Result: The system displays a Process Member menu. On the Process Member menu, type the number of the option DO NOW and press Enter. Result: The system executes the statement in the Accum member.

Introduction

This chapter explains prerequisites, options, and procedures for creating a database. The chapter “Introduction to Creating Objects” on page 209 is a prerequisite for this chapter.

Chapter contents

Overview of Chapter	236
Access to the Create Database Panel	237
Prerequisites for Creating a Database	238
Beginning to Create the Database	239
Using Values from an Existing Database	240
Specifying an Existing Storage Group for the New Database	242
Creating a New Storage Group for the New Database	244

Overview of Chapter

Prerequisite for using this chapter

Before following the instructions in this chapter, review the chapter “Introduction to Creating Objects” on page 209.

Organization of information in the chapter

The chapter provides background information about accessing the Create Database panel, prerequisites for creating a database (regardless of the method you use from the Create Database panel), instructions for beginning the process of creating a database, and instructions for the three creation methods possible from the panel.

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Access	Prerequisites	Beginning	Model	Existing Storage Group	New Storage Group
	p.237	p.238	p.239	p.240	p.242	p.244
Review methods of access (and the method to exit)	√					
Create a new database using an existing database as a model		√	√	√		
Create a new database and use an existing storage group		√	√		√	
Create a new database and a new storage group		√	√			√

Access to the Create Database Panel

Overview

The unit lists methods that result in display of the Create Database panel and explains how to exit the panel if you do not want to create a new database at this time.

Methods to access the Create Database panel

The system displays the Create Database panel in these cases.

- You accessed the Create Database panel from the Create Menu.
- You are creating another object that requires the entry of the name of a database and you enter a name of an object that is not defined in the DB2 catalog or is not currently stored in memory that is active during a session when you use the Create Menu.

For a reminder about instructions for the methods to access the Create Menu (including the fast path), see “Introduction to Creating Objects” on page 209.

If you do not want to create a database at this time

If you do not want to create a database at this time, follow these steps.

Step	Action
1	On the Create Database panel, blank out the new name in the Create Database field.
2	Type END on the command line and press Enter, or press the appropriate function key.

Prerequisites for Creating a Database

Overview

The unit contains the prerequisites required to create a database.

Authorizations required to create a table

To create a table, you must have one of these levels of authority:

- SYSADM authority
 - CREATEDBA privilege (With this level of privilege, the system automatically grants you DBADM authority over the database.)
 - CREATEDBC privilege (With this level of privilege, the system automatically grants you DBCTRL authority over the database.)
-

General prerequisites for creating a database

You must meet one of these general prerequisites for creating a database:

- An existing storage group
- A storage group which you create during the process of creating the database

Beginning to Create the Database

Overview

The unit contains the instructions applicable to the three options available when creating a database.

Procedure to follow to begin to create the database

Follow these steps to begin to create the database.

Step	Action
1	On the Create Database panel, type the name of the new database in the Create Database field.
2	Press Enter. Result: The system displays a message directing you to select a database to use as a model and moves the cursor to the Select field for the first database in the list.

Determining what to do next

The chart lists the three options !DB/WORKBENCH provides from the Create Database panel to create a database.

Characteristics of the New Database	Unit to Use
A new database that uses the same storage group and buffer pool as an existing database	“Using Values from an Existing Database” on page 240
A new database that uses a storage group that is defined in the DB2 catalog	“Specifying an Existing Storage Group for the New Database” on page 242
A new database that uses a new storage group that you want to create at this time	“Creating a New Storage Group for the New Database” on page 244

Using Values from an Existing Database

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Database” on page 239 before you proceed with this unit, which covers one of the three options available for creating a database.

This unit covers the information needed if you want to use the values from an existing database as the source of values for a new database. It describes the prerequisites and the procedures for this option, and gives instructions on what to do after you create the database.

Prerequisites for using values from an existing database

To perform the procedures in this unit, you need:

- The name of the existing database you want to use as a model (as the source of values for the storage group and the buffer pool) for the new database
 - The setting for the ROSHARE parameter (if desired)
-

Procedure for using values from an existing database

To use values from an existing database, follow these steps.

Step	Action
1	On the Create Database panel, type S in the Select field for the database whose values you want to use for the new database.
2	Press Enter. Result: The system records and displays values from the model database in the Default Storage Group field and in the Default Buffer Pool field.
3	Press Enter. Result: The system stores in memory the CREATE statements using values for the database name, storage group, buffer pool, and ROSHARE parameter (if entered), and redisplay the Create Menu.

Determining what to do next

You have two broad choices from the Create Menu:

- If you want to create other objects, type the letter of the option in the Option field and press Enter.
- If you want to perform a task, such as to process the statements or save the statements from memory to a member in the user PDS, type **END** on the command line and press Enter (or press the appropriate function key) to display the Process Member menu for the Create Menu.

If you need a reminder about using the Process Member menu from the Create Menu, see the chapter “Introduction to Creating Objects” on page 209.

Specifying an Existing Storage Group for the New Database

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Database” on page 239 before you proceed with this unit, which contains one of the three options available for creating a database.

This unit covers the information needed if you want to use a specific storage group (and its associated buffer pool) for the new database you are creating. The unit describes the prerequisites and the procedure for this option, and gives instructions on what to do after you create the database.

Prerequisites for specifying a storage group

To perform the procedure in this unit, you need:

- The name of the existing storage group you want to use with the new database
 - The setting for the ROSHARE parameter (if desired)
-

Procedure for specifying an existing storage group

To specify an existing storage group, follow these steps.

Step	Action
1	On the Create Database panel, type ? in the Default Storage Group field.
2	Press Enter. Result: The system displays a list of storage groups.
3	Type S in the Select field for the storage group you want to specify for the database.
4	If you want to provide a value for ROSHARE, type the value in the ROSHARE field.
5	Press Enter. Result: The system stores in memory the CREATE statements using the values for the database name, the storage group, and the buffer pool, and redisplay the Create Menu.

Determining what to do next

You have two broad choices from the Create Menu:

- If you want to create other objects, type the letter of the option in the Option field and press Enter.
- If you want to perform a task, such as to process the statements or save the statements from memory to a member in the user PDS, type **END** on the command line and press Enter (or press the appropriate function key) to display the Process Member menu for the Create Menu.

If you need a reminder about using the Process Member menu from the Create Menu, see the chapter “Introduction to Creating Objects” on page 209.

Creating a New Storage Group for the New Database

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Database” on page 239 before you proceed with this unit, which covers one of the three options available for creating a database.

This unit covers information needed if you want to create a new storage group as well as a new database. This unit describes the three parts to the procedure for the option, and gives instructions on what to do after you create the database.

Prerequisites for creating a new storage group

To perform the procedure in this unit, you need:

- The setting for the ROSHARE parameter (if desired)
- The prerequisites for the option you select from those available using the Create Storage Group panel

For information about prerequisites for creating a storage group, see the chapter “Creating a Storage Group” on page 305.

Beginning the procedure

Follow these steps to begin the procedure to create a new storage group by accessing the Create Storage Group panel.

Step	Action
1	On the Create Database panel, type the name of the new storage group in the Default Storage Group field.
2	Press Enter. Result: The system displays the Create Storage Group panel.

Creating the storage group

To create the new storage group, follow the instructions in the chapter “Creating a Storage Group” on page 305.

Completing the procedure

When you have completed the creation of the new storage group for the new database, the system redisplay the Create Database panel. Follow these steps.

Step	Action
1	On the Create Database panel, confirm that the values in the fields for storage groups and buffer pool are correct.
2	If you want to provide the value for the ROSHARE parameter, type the value in the ROSHARE field.
3	Press Enter. Result: The system stores in memory the CREATE statements for the new database using the values for the database name, the storage group, and buffer pool, and the CREATE statements for the storage group. The system also redisplay the Create Menu.

Determining what to do next

You have two broad choices from the Create Menu:

- If you want to create other objects, type the letter of the option in the Option field and press Enter.
- If you want to perform a task, such as to process the statements or save the statements from memory to a member in the user PDS, type **END** on the command line and press Enter (or press the appropriate function key) to display the Process Member menu for the Create Menu.

If you need a reminder about using the Process Member menu from the Create Menu, see the chapter “Introduction to Creating Objects” on page 209.

Creating a New Storage Group for the New Database.

Introduction

This chapter explains prerequisites, options, and procedures for creating a table space (whether simple, segmented, or partitioned). The chapter “Introduction to Creating Objects” on page 209 is a prerequisite for this chapter.

Chapter contents

Overview of Chapter	248
Access to the Create Table Space Panel	250
Common Prerequisites for Creating a Table Space	251
Specific Prerequisites and Tasks for a Segmented Table Space	253
Specific Prerequisites and Tasks for a Partitioned Table Space	254
Beginning to Create the Table Space	256
Using Values from an Existing Table Space	257
Using the Default Values for the New Table Space	260
Completing a Non-Partitioned, DB2-Managed Dataset	263

Overview of Chapter

Prerequisite for using this chapter

Before following the instructions in this chapter, review the information in “Introduction to Creating Objects” on page 209.

Organization of information in the chapter

The chapter provides background information about accessing the Create Table Space panel, prerequisites for creating a table space (whether simple, segmented, or partitioned), and instructions for creating a table space. It also includes an introduction to setting up datasets for the table space.

This chapter and instructions common to all types of table spaces

Most of the instructions in this chapter are applicable to any type of table spaces. When additional instructions specific to segmented or partitioned table spaces are needed, the information will always be distinguished from general instructions for all types of table spaces. For example, if you are creating a simple or segmented table space, you can skip a unit with a heading such as “Specific Prerequisites and Tasks for a Segmented Table Space” on page 253.

This chapter and instructions for segmented or partitioned table spaces

This chapter contains instructions and prerequisites for creating a segmented or partitioned table space. The instructions specific to segmented or partitioned table spaces will always be distinguished from general instructions that apply to all types of table spaces. For example, they will have a distinctive heading such as the one for this paragraph.

The chapter and methods for preparing VSAM datasets

This chapter includes an introduction to the three methods available within !DB/WORKBENCH for you to set up the VSAM dataset for the table space. The chapter provides instructions for one of these methods: for creation of single or segmented table spaces with a DB2-created and managed dataset. (If you are using one of the other methods to set up the VSAM dataset for the table space or are setting up a partitioned table space, the instructions are provided in the chapter “Completing Requirements for VSAM Datasets” on page 315.)

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Access	Common Prerequisites	Segmented Prerequisites	Partitioned Prerequisites	Beginning	Model	Defaults	Non-partitioned
	p.250	p.251	p.253	p.254	p.256	p.257	p.260	p.263
Review methods of access (and the method to exit)	√							
Create a new simple table space using an existing table space as a model		√			√	√		
Create a new simple table space using default values		√			√		√	
Create a new segmented table space using an existing table space as a model		√	√		√	√		
Create a new segmented table space using default values		√	√		√		√	
Create a new partitioned table space using an existing table space as a model		√		√	√	√		
Create a new partitioned table space using default values		√		√	√		√	
Completing requirements for a simple or segmented table space with a VSAM dataset that DB2 creates and manages								√

Access to the Create Table Space Panel

Overview

The unit lists methods that result in display of the Create Table Space panel and explains how to exit the panel if you do not want to create a new table space at this time.

Methods to access the Create Table Space panel

The system displays the Create Table Space panel when:

- You access the Create Table Space panel from the Create Menu
- You are creating another object that requires the entry of the name of a table space and you enter a name of an object that is not defined in the DB2 Catalog or is not currently stored in memory that is active during a session when you use the Create Menu

For a reminder about instructions for the methods to access the Create Menu (including the fast path), see “Introduction to Creating Objects” on page 209.

If you do not want to create a table space at this time

If you do not want to create a table space at this time, follow these steps.

Step	Action
1	On the Create Table Space panel, blank out the new name in the Create Table Space field.
2	Type END on the command line and press Enter, or press the appropriate function key.

Common Prerequisites for Creating a Table Space

Overview

The unit contains the prerequisites required to create a table space (whether simple, segmented, or partitioned).

Terminology

With DB2 and !DB/WORKBENCH, a table space is the area of a database that contains tables. !DB/WORKBENCH supports simple, segmented, or partitioned table spaces.

With DB2 and !DB/WORKBENCH, each table space can consist of one or more VSAM datasets. !DB/WORKBENCH supports these VSAM datasets:

- A VSAM dataset that DB2 creates and manages
- A VSAM dataset that you have previously created and that you manage
- A VSAM dataset that you create when you create the table space

For information about which type of table space and the type of management of the VSAM dataset is best for your needs, see the *IBM DB2 Administration Guide*.

Authorizations required to create a table space

To create a table space, you must have one of these levels of authority:

- SYSADM authority
- DBADM authority
- DBCTRL authority
- DBMAINT authority
- CREATETS privilege

General prerequisites for creating a table space

You must have a VSAM dataset previously defined, or you must define one during these procedures. You also must meet one of these general prerequisites for creating a table space:

- An existing database
- A database which you create during the process of creating the table space

Prerequisites common to simple, segmented, or partitioned table spaces

Whether you use or modify the values from an existing table spaces or from the defaults with !DB/WORKBENCH or whether you enter all new values, you must enter or confirm these values:

- Buffer pool name
- CLOSE setting
- LOCKsize setting
- SEGsize setting (The default setting of 0 results in a table space that is either simple or partitioned.)
- Partition value (The default setting of 0 or 1 results in a table space that is either simple or segmented.)
- COMPRESS setting (You specify a percentage between 0–99 for compression, for example, 50 percent.)

Prerequisites for the VSAM dataset

When you create the table space, you also must set up values related to the VSAM dataset. You enter or confirm these values:

- ERASE setting
- Priqty value
- Secqty value
- PCTfree value
- FREE page value
- Gbpcache value
- VSAM password
- Storage group name

Specific Prerequisites and Tasks for a Segmented Table Space

Overview

The unit contains a general overview of tasks and prerequisites required for creating segmented table spaces.

General overview of creating a segmented table space

With !DB/WORKBENCH, you perform only one task to generate statements to create a segmented table space. When you create the table space, you enter a value for SEGsize (other than 0).

Prerequisites specific to segmented table spaces

In addition to the prerequisites common to all types of table spaces, you also need the value you want for SEGsize.

Specific Prerequisites and Tasks for a Partitioned Table Space

Overview

The unit contains background information, a general overview of tasks, and prerequisites required for creating partitioned table spaces.

Background about partitioned table spaces

In a partitioned table space, the available space is divided into separate units of storage called partitions, each containing a part of a table. Each partition can be assigned to a separate storage group; however, the entire collection of data is logically a single table.

A partitioned table space can have only one table and can consist of up to 64 partitions. The maximum size of any table space is 64 gigabytes, but the maximum size of any partition is 4 gigabytes. Therefore, a table space with fewer than 16 partitions cannot reach maximum size.

General overview of creating a partitioned table space

With !DB/WORKBENCH, you perform these tasks to generate statements to create a partitioned table space. First, when you create the table space, you enter a value for Partitions (other than 0 or 1). Second, when you set up the method you will use to create and manage the VSAM dataset, you perform one of two tasks:

- You specify that all partitions have the same value. (No other action is needed.)
- You specify that all partitions have different values, and complete panels specifying the values for each partition.

Creating other partitioned objects

When you complete creation of the table space and set up of VSAM dataset (whether DB2 manages them or you do), you should create the other partitioned objects in the order shown on the Create Menu. For example, next you should create the tables and then the indexes.

Prerequisites specific to partitioned table spaces

In addition to the prerequisites common to all types of table spaces, you need to determine the number of partitions the table space needs. You also need to determine whether all partitions will have the same values, and, if they do not, what values will apply to each partition.

Beginning to Create the Table Space

Overview

The unit contains the instructions applicable to all options available when creating a table space.

Procedure to follow to begin to create the table space

Follow these steps to begin to create the table space.

Step	Action
1	On the Create Table Space panel, perform the appropriate action: <ul style="list-style-type: none"> ● Type the name of the database in the Database field. ● Type S in the Select field for the database you want to associate with the new table space.
2	Press Enter. Result: The system displays a message directing you to select a table space to use as a model and moves the cursor to the Select field for the first table space in the list.

Determining what to do next

The chart lists the options !DB/WORKBENCH provides from the Create Table Space panel to create a table space.

Characteristics of the New Table Space	Unit to Use
A new table space that uses the same values as an existing table space	“Using Values from an Existing Table Space” on page 257
A new table space that uses default values (either from the system or from the default options you set up)	“Using the Default Values for the New Table Space” on page 260

Using Values from an Existing Table Space

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Table Space” on page 256 before you proceed with this unit, which covers one of the options available for creating a table space.

This unit covers the information needed if you want to use the values from an existing table space as the source of values for a new table space. It describes the prerequisites and the procedures for this option, and gives instructions on what to do after you create the table space.

Prerequisites for using values from an existing table space

To perform the procedures in this unit, you need:

- The name of the existing table space you want to use as a model (as the source of such values as the CLOSE setting and the LOCKsize) for the new table space
 - The setting for the SEGsize (if you are creating a segmented table space)
 - The setting for Partitions (if you are creating a partitioned table space)
-

Background about restricting the display of table spaces

If you want to restrict the display of table spaces to those in the database, follow these steps. (Typing an asterisk over the first character in the table space name does not affect the name.)

Step	Action
1	On the Create Table Spaces panel, type * (asterisk) in the Create Tables field.
2	Press Enter. Result: The system redisplayes the new table space name (and replaces the asterisk) and lists only those table spaces in the database that you specified.

Reminder about segmented and partitioned table spaces

If you are planning to create a segmented or partitioned table space, be sure to modify these values at this time.

- SEGsize (for segmented table spaces)
- Partition (for partitioned table spaces)

Procedure by specifying an existing table space

To specify an existing table space, follow these steps.

Step	Action
1	On the Create Table Space panel, type S in the Select field for the table space whose values you want to use for the new table space.
2	Press Enter. Result: The system displays a panel showing the completed values for buffer pool, CLOSE, and LOCKsize (as well as values needed for segmented or partitioned table spaces) using the values in the model table space you selected.
3	If needed, modify the values.
4	Press Enter. Result: The system displays a panel showing the alternatives for creating and managing the VSAM dataset.

Determining what to do next

Review the chart to determine what to do next.

IF you created...	AND you want...	THEN use...
a simple or segmented table space	DB2 to create and manage your VSAM dataset	the instruction in the unit “Completing a Non-Partitioned, DB2-Managed Dataset” on page 263
a partitioned table space	DB2 to create and manage your VSAM dataset	the chapter :href refid=cvsam page=yes.
a simple, segmented, or partitioned table space	to manage the VSAM dataset which you previously created	the chapter :href refid=cvsam page=yes.
a simple, segmented, or partitioned table space	to create a VSAM dataset at this time and to manage the dataset yourself	the chapter :href refid=cvsam page=yes.

Using the Default Values for the New Table Space

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Database” on page 239 before you proceed with this unit, which contains one of the three options available for creating a table space.

This unit covers the information needed if you want to use default values from !DB/WORKBENCH or the default values you established using the default option (O on the Create Menu) for the new table space you are creating. The unit describes the prerequisites and the procedure for this option, and gives instructions on what to do after you create the table space.

Reminder about default values

If you need a reminder about how the system uses default values in this process, see “Introduction to Creating Objects” on page 209.

Reminder about segmented and partitioned table spaces

If you are planning to create a segmented or partitioned table space, be sure to modify these values at this time.

- SEGsize (for segmented table spaces)
- Partition (for partitioned table spaces)

Procedure for using default values for the new table space

To use default values for the new table space, follow these steps.

Step	Action
1	On the Create Table Space panel, type the name of the new table space in the Create Table Spaces field.
2	Press Enter. Result: The system displays a list of table spaces available as a model.
3	Press Enter again.
4	Press Enter again. Result: The system displays a panel showing the completed values for Buffer pool, CLOSE, and LOCKsize (as well as the values needed for segmented or partitioned table spaces) using the default values.
5	If needed, modify values.
6	Press Enter. Result: The system displays a panel showing all values for creating and managing the VSAM dataset.

Determining what to do next

Review the chart to determine what to do next.

IF you created...	AND you want...	THEN use...
a simple or segmented table space	DB2 to create and manage your VSAM dataset	the instruction in the unit “Completing a Non-Partitioned, DB2-Managed Dataset” on page 263
a partitioned table space	DB2 to create and manage your VSAM dataset	the chapter “Completing Requirements for VSAM Datasets” on page 315
a simple, segmented, or partitioned table space	to manage the VSAM dataset which you previously created	the chapter “Completing Requirements for VSAM Datasets” on page 315
a simple, segmented, or partitioned table space	to create a VSAM dataset at this time and to manage the dataset yourself	the chapter “Completing Requirements for VSAM Datasets” on page 315

Completing a Non-Partitioned, DB2-Managed Dataset

Overview

This unit contains instructions only if you have created a simple or segmented table space and if you want DB2 to create and manage the VSAM dataset for you. This unit contains a reminder and the procedure to follow.

Reminder about the Space Estimator

You can use the !DB/WORKBENCH Space Estimator to estimate the value for Priqty. For an overview of the procedure to use the estimator, see the unit “Methods Available to Estimate Requirements for Space” on page 222. For more detailed information, see “Using the DASD Estimators” on page 111.

Procedure to prepare a DB2-created and managed VSAM dataset

Follow these steps to set up a DB2-created and managed VSAM dataset.

Step	Action
1	On the panel listing alternatives for tables spaces, type 1 in the Choose One field.
2	Press Enter. Result: The system displays a panel listing the values to be used for the VSAM dataset for the new table space.
3	If necessary, modify the values.
4	If the storage group name is not displayed, perform the appropriate action: <ul style="list-style-type: none"> ● Type the name of the storage group in the Storage Group field. ● Type S in the Select field for the storage group you want to use.
5	Press Enter. Result: The system displays a panel that allows you to create a new table within the table space.

Determining what to do next

The chart lists your choices and the appropriate action and source.

IF you...	THEN...
want to create a table	type Y and press Enter. Result: The system displays the Create Tables panel. For more information, see the chapter “Creating a Table” on page 265.
do not want to create a table	type N and press Enter. Result: The system redisplay the Create Menu.

Introduction

This chapter explains prerequisites, background, and procedures for creating a table. The chapter “Introduction to Creating Objects” on page 209 is a prerequisite for this chapter.

Chapter contents

Overview of Chapter	266
Access to the Create Table Panel	267
Prerequisites for Creating a Table	268
Beginning to Create the Table	269
Providing Required Values for the New Table	270
Background about Functions and Selects Available	272
Background about Using the Editor	274
Example of How to Use the Editor	276
Defining the Columns for the New Table	278
Adding Comments and Labels to Columns	280
Completing Values for the Tables	281
Creating Synonyms for the Table	282

Overview of Chapter

Prerequisite for using this chapter

Before following the instructions in this chapter, review the information in “Introduction to Creating Objects” on page 209

Organization of information in the chapter

The chapter provides background about accessing the Create Table panel and prerequisites for creating a table. It discusses the use of optional synonyms and provides background about functions and selects available when defining columns. It discusses using the editor and provides instructions for creating the table.

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Access	Prerequisites	Beginning	Required Values	Functions and Selects	The Editor	Example	Defining Columns	Column Comments	Other Values	Synonyms
	p.267	p.268	p.269	p.270	p.272	p.274	p.276	p.278	p.280	p.281	p.282
Review methods of access (and the method to exit)	√										
Perform preparatory steps if you accessed the Create Tables panel from the Create Menu		√									
Create a new table for the first time		√		√	√	√	√	√	√	√	√
Create a new table after you have learned about functions, selects, and the editor		√		√				√	√	√	

Access to the Create Table Panel

Overview

The unit lists methods that result in display of the Create Table panel and explains how to exit the panel if you do not want to create a new table at this time.

Methods to access the Create Table panel

The system displays the Create Table panel in these cases.

- You accessed the Create Table panel from the Create Menu.
- You are creating another object that requires the entry of a table and you enter a name of an object that is not defined in the DB2 catalog or is not currently stored in memory that is active during a session when you use the Create Menu.

For a reminder about instructions for the methods to access the Create Menu (including the fast path), see “Introduction to Creating Objects” on page 209.

If you do not want to create a table at this time

If you do not want to create a table at this time, follow these steps.

Step	Action
1	On the Create Table panel, blank out the new name in the Create Table field.
2	Type END on the command line and press Enter, or press the appropriate function key.

Prerequisites for Creating a Table

Overview

The unit contains the prerequisites required to create a table.

Terminology

All data in a DB2 database is presented in tables. A table is a collection of rows that all have the same columns. When creating a table in DB2, you are defining an ordered set of columns. A column contains all occurrences of one of the entities of a table. At the intersection of a column and row is a value. The storage representation of a row is a record.

Authorizations required to create a table

To create a table, you must have one of these levels of authority:

- SYSADM authority
 - CREATETAB privilege for the database
 - DBADM authority
 - DBCTRL authority
 - DBMAINT authority
-

General prerequisites for creating a table

The table is owned by a database and table space; therefore, you must provide the name of these objects when creating a table:

- An existing database or one which you created in this session using options on the Create Menu
 - An existing table space or one which you create in this session using options on the Create Menu
-

Reminder

When you provide the name of the table space in this process and the one for creating a table space, you can type an * (asterisk) in the Table Space field and press Enter. The system lists only those table spaces in the database you specified.

Beginning to Create the Table

Overview

The unit contains the instructions needed if you began to create a table by accessing the Tables option on the Create Menu. (These steps are not necessary if you are creating a table for a new database or table space because the program will carry the values for the database or table space from your previous entry.)

Procedure to follow to begin to create the table

Follow these steps to begin to create the table.

Step	Action
1	On the Create Table panel perform the appropriate action: <ul style="list-style-type: none"> ● Type the name of the database in the Database field. ● Type S in the Select field for the database you want.
2	Press Enter. Result: The system displays a message directing you to type the name of the table space.
3	Perform the appropriate action: <ul style="list-style-type: none"> ● Type the name of the table space in the Table Space Name field. ● Type S in the select field for the table space you want.
4	Press Enter. Result: The system displays a panel listing the values you provided for the database and the table space. The system also displays your USERID in the Creator ID field.

Providing Required Values for the New Table

Overview

If you began to create the table by accessing the tables option on the Create Menu, be sure you have followed the instructions in the unit “Beginning to Create the Table” on page 269 before you proceed with this unit.

This unit covers the prerequisites and the procedures for providing the required values to create a new table.

Prerequisites for providing required values for a new table

To perform the procedures in this unit, you need:

- The name for the new table (The name must be one that is different from any table, view, or synonym associated with the Creator ID.)
 - The values required to define the columns for the new table or the table(s) whose column definitions you plan to copy for column definitions in the new table. (You use these values in the unit “Defining the Columns for the New Table” on page 278.)
-

Method to avoid duplicate names for tables

You can display a list of all tables associated with the Creator ID by typing an * (asterisk) in the Create Table field. The system displays the tables for the specified Creator ID. (The tables do not serve as models for the table you are creating, but as reminders of the names you have used previously.)

Procedure for providing values to identify the table

Follow these steps.

Step	Action
1	If you want to create the table under an ID different from your own, type the creator ID in the Creator ID field.
2	Type the name of the new table in the Create Table field.
3	Press Enter. Result: The system displays a panel that provides instructions and functions and selects for defining columns for the new table.

Background about Functions and Selects Available

Overview

This unit provides background about functions and selects available on the panel for defining columns.

Overview of tasks you can perform on the panel

You can perform two broad types of tasks on the panel for defining columns for the new table.

Task You Want to Perform	Action to Use the Option
Work with the column definitions for the table you are creating	Type the letter of the function in the Processing on your column's field and press Enter.
Use a column definition from another table	Type the letter of the select in the Select field for the table and press Enter.

Functions for working with the columns you are defining

Review the chart for the functions you can use with the columns you are defining. You can reuse these functions as many times as you need.

Task You Want to Perform	Action to Use the Option
Display the column's definition for the new table (You can scroll to view the columns.)	L
Invoke the ISPF Editor so you can edit the definitions	E

Selects for using column definitions from other tables

Review the chart for the selects available using column definitions from other tables.

Task You Want to Perform	Action to Use the Option
Append the column definitions in the selected table to the column definitions for the new table. (You can use this select with several tables at one time.)	A
Display the table columns (which lists panel values such as columns, type, and length) for the selected table. (If desired, you can use any command or select on the Table Columns panel.)	L
Replace the existing column definition for the new table with the column definition in the selected table.	R

Reminder about the Table Columns panel

The system displays the Table Columns panel when you use the C (Columns display) select with a table on the Tables panel. You also see the panel when creating a table.

The data shown on the Table Columns panel is the same data shown if you use the L function in the processing on your column field. The difference is that the Table Columns panel displays data about column definitions from an existing table, and the other panel displays data from definitions you are creating.

Background about Using the Editor

Overview

This unit contains background about using the Editor when you define columns.

Background about using the Editor

When you type the E function in the processing on your column's field and press Enter, the system invokes the ISPF Editor and displays a panel with instructions.

Functions performed by the Editor with SAVE

The Editor performs these operations when you type SAVE on the command line and press Enter.

- Save the column definition you have added, edited, or replaced during the session
- Check the input you have made and generated a message indicating problems or omissions. (The message for a column definition will be displayed above the line for the definition.)
- Completes the name of the data type (or F/D TYPE) if needed (For example, if you typed C, the system displays CHAR.)
- Replaces any combination of a NO value listed in both the Nulls? column and the Defaults? columns with a NO in the Nulls column only. (When the system generates the statement, it uses NOT NULL WITH DEFAULT in these cases.)
- Adds a sequence number for duplicate columns. (For example, two columns named DEPTTYPE become DEPTTYPE and DEPTTYPE_02.)

Commands available in the Editor

You can use these commands within the Editor to perform tasks.

Task You Want to Perform	Command to Use
Saves all definitions and processes the input	SAVE
Cancels all edits made during the session	CANCEL

General operations in the ISPF Editor

The operations in the Editor follow ISPF practices with these exceptions:

- Use the TAB to move from field to field.
- Type only the first character for the type of data. (For example, type C if the column will be for character data.)

Choices available for errors in definition

If your column definitions result in an error message, you have two choices:

- Use the CANCEL command to end the session without saving the column definition.
- Correct the definition and use the SAVE command again.

Example of How to Use the Editor

Overview

Before you review the example, be sure you have reviewed the units “Providing Required Values for the New Table” on page 270 and “Background about Using the Editor” on page 274. This unit contains a description of the characteristics of the example.

Characteristics of the example of how to use the editor

In the example, you have previously provided the values for owning database and table space and for the new name of the table. The system has displayed a panel listing the functions (Edit and List) and the selects (A, L, and R) you can use on the panel. These characteristics exist for the example shown in this unit:

- To avoid keying, you plan to use two existing tables that contain similar column definitions to the ones you need.
 - You plan to change the names for five columns
 - You plan to add one column definition
 - You forget to type the value for FLDLEN or (length) for the new column definition.
-

Copying the column definitions for the example

In this part of the example of the procedure, you copy column definitions from the two tables you plan to use as a source and invoke the editor.

Step	Action
1	Type A in the Select field for the table.
2	Type A in the Select field for the second table.
3	Press Enter. Result: The system copies the column definitions from the first table and then adds the definitions from the second table. The system moves the cursor to the Processing on your columns field.
4	Type E in the Processing on your list field.
5	Press Enter. Result: The system invokes the ISPF editor and displays the column definition for the two tables.

Editing and saving the column definition for the example

In this portion of the example, you edit the definitions, and save the definition.

Step	Action
1	Type over the names of the five column names.
2	Insert the line where you want the additional column definition.
3	Type the new column definition.
4	Type SAVE .
5	Press Enter. Result: The system processes the column statements and issues a message in the line above the column definition indicating you need to provide the value for FLDLEN for the new column definition.

Correcting the error and completing the column definition for the example

In this portion of the example, you correct the error and save the definition.

Step	Action
1	Type the value for FLDLEN in the column on the panel.
2	Type SAVE .
3	Press Enter. Result: The system redisplay the panel listing functions and selects. (No values are in the Processing on your columns on the select fields for the tables.)
4	Press Enter. Result: The system displays a panel with optional fields for a comment line and a label for the column field.

Defining the Columns for the New Table

Overview

This unit provides prerequisites and instructions for defining the columns for the new table. (If you have not previously defined columns using the panel and the editor, review these units “Background about Functions and Selects Available” on page 272, “Background about Using the Editor” on page 274, and “Example of How to Use the Editor” on page 276.)

Prerequisites for defining a column

To define a column, you provide these values:

- The name of the column
- The type of data in the field (such as integer or character data)
- The length of the field
- The scale (if you use a decimal data type)
- Whether nulls data (or missing data) is allowed
- Whether you want a not-null column to have a default value

Process for defining a column

Follow these steps.

Step	Action
1	If appropriate for the columns definition for the new table, use the selects on the panel to copy column definitions to be used by the editor.
2	Type E in the Processing on your columns field.
3	Press Enter. Result: The system invokes the ISPF Editor and displays the contents of the editor.
4	Edit the column definitions.
5	Type SAVE on the command line
6	Press Enter. Result: The system displays a panel listing labels and comment fields for each column you placed in the file.

Adding Comments and Labels to Columns

Overview

This unit provides instructions for adding optional comments and labels for columns.

Prerequisites

To perform this procedure you need to know the label and comment you want to use for each column definition. You can, however, choose to not provide labels or comments.

Adding labels or comments to the columns

If you want labels or comments for the column, type them during this procedure. (If you want a comment or label for the table itself, do not type them on this panel.) Follow these steps.

Step	Action
1	On the Columns Detail List, type the text for the labels or comments for the column definitions you want.
2	Press Enter. Results: The system display fields for entry of values to restrict the table.

Completing Values for the Tables

Overview

This unit lists values for the table that are required and those that are optional and provide restrictions.

If your database uses the ROSHARE parameter

If your database uses the ROSHARE parameter, you must provide the object ID (the value for OBID).

If you want additional values

You can provide these additional values to the system.

- Editproc name
 - Validproc name
 - Audit status
 - Data capture status
 - Restrict w/Drop
 - Label for the table
 - Comment on the table
-

Procedure for adding restrictions

Follow these steps.

Step	Action
1	If your database uses the ROSHARE parameter, type the object ID in the OBID field.
2	If you want to use any of the restrictions, type the value in the field.
3	Press Enter. Results: The system displays a panel asking you if you want to create an index.

Creating Synonyms for the Table

Overview

Use this unit if you want to create synonyms for the table.

Value used in the procedure

You can provide a maximum of eight synonyms for the table.

Procedure for creating synonyms for the table

If you want synonyms for the table follow these steps. (If you have created a partitioned table space, you should also create a partitioned index.)

Step	Action
1	On the Create Table panel, type each synonym in the Synonym field.
2	Press Enter. Results: The system displays a panel asking if you want to create an index for the new table.

Determining what to do next

The chart lists your choices and the appropriate action and source.

IF you...	THEN...
want to create an index	type Y and press Enter. Result: The system displays the Create Index panel. For more information, see “Creating an Index and Index Space” on page 283
do not want to create an index	type N and press Enter. Result: The system redisplay the Create Menu.

Chapter 12. Creating an Index and Index Space

Introduction

This chapter explains prerequisites, options, and procedures for creating an index and index space (whether simple or partitioned). The chapter “Introduction to Creating Objects” on page 209 is a prerequisite for this chapter.

Chapter contents

Overview of Chapter	284
Access to the Create Index Panel	286
Common Prerequisites for Creating an Index and Index Space	287
Specific Tasks for Unique or Clustering Indexes	290
Specific Tasks for Partitioned Index and Index Space	291
Beginning to Create the Index	292
Using Values from an Existing Index	294
Using the Default Values for the New Index	296
Selecting Table Columns to Be Indexed	298
Completing a Non-Partitioned, DB2-Managed Dataset	302

Overview of Chapter

Prerequisite for using this chapter

Before following the instructions in this chapter, review the information in “Introduction to Creating Objects” on page 209.

Organization of information in the chapter

The chapter provides background information about accessing the Create Index panel, prerequisites for creating an index (whether unique, clustering, or partitioned) and index spaces (whether simple or partitioned), and instructions for creating an index. It also includes an introduction to setting up data sets for the index space.

This chapter and instructions common to all types of indexes

Most of the instructions in this chapter are applicable to any type of index. When additional instructions specific to a unique, clustering or partitioned index are needed, the information will always be distinguished from general instructions for all types of indexes. For example, if you are creating a simple index space, you can skip a unit with a heading such as “Additional Procedures for Partitioned Index Spaces.”

This chapter and instructions for partitioned indexes and index spaces

This chapter contains instructions and prerequisites for creating a partitioned index or index space. The instructions specific to partitioned index space will always be distinguished from general instructions that apply to all types of index spaces. For example, they will have a distinctive heading such as the one for this paragraph.

The chapter and methods for preparing VSAM datasets

This chapter includes an introduction to the three methods available within !DB/WORKBENCH for you to set up the VSAM dataset for the index space. The chapter provides instructions for one of these methods: for creation of single index spaces with a DB2-created and managed dataset. (If you are using one of the other methods to set up the VSAM dataset for the index space or are setting up a partitioned index space, the instructions are provided in the chapter “Completing Requirements for VSAM Datasets” on page 315.)

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Access	Common Prerequisites	Unique or Clustering	Partitioned	Beginning	Model	Defaults	Table Columns	Non-partitioned
	p.286	p.287	p.290	p.291	p.292	p.294	p.296	p.298	p.302
Review methods of access	√								
Create a new unique or clustering index using an existing index as a model		√	√		√	√		√	
Create a new unique or clustering index using default values		√	√		√		√	√	
Create a new partitioned index using an existing index as a model		√	√	√	√	√		√	
Create a new partitioned index using default values		√	√	√	√		√	√	
Complete requirements for a unique or clustering index with a VSAM dataset that DB2 creates and manages									√

Access to the Create Index Panel

Overview

The unit lists methods that result in display of the Create Indexes panel and explains how to exit the panel if you do not want to create a new index at this time.

Methods to access the Create Index panel

The system displays the Create Index panel in these cases.

- You accessed the Create Index panel from the Create Menu.
- You enter yes in response to a message asking you if you want to create an index.

For a reminder about instructions for the methods to access the Create Menu (including the fast path), see “Introduction to Creating Objects” on page 209.

Common Prerequisites for Creating an Index and Index Space

Overview

The unit contains background and the prerequisites required to create an index (whether unique, clustering, or partitioned) and index spaces (whether simple or partitioned).

Terminology

With DB2 and !DB/WORKBENCH, an index is an ordered set of pointers to the data in a DB2 table. Each index is based on the values of data in one or more columns of a table; and you can create more than one index for a table. You can use indexes for two broad purposes:

- To attempt to improve performance
- To insure uniqueness

Except for changes in performance, users of a table are unaware that an index is being used. The usefulness of an index depends on the index key, which is a single column or an ordered collection of columns on which an index is defined. Columns that are used frequently to perform selections, joins, projections, counting, groupings, and ordering operations are good candidates for keys.

Types of indexes

There are three types of indexes.

- Unique** An index is unique if the table does not contain two or more rows with the same value as the index. In other words, there are no duplicates. The system enforces uniqueness when you update or add new rows.
- Clustering** A clustering index determines the approximate order in which the system stores records in a table; DB2 can access the entire table in the sequence of the clustering key faster than any other sequence. Each table can have only one clustering index.
- Partitioned** When you create a clustering index on a table in a partitioned table space, the index is also partitioned into multiple index spaces.

Maintenance and use of indexes

After you create an index, DB2 maintains it and determines whether or not to use the index to access the table. You can, however, check, repair, or reorganize the index.

Timing of creation of indexes

You can create an index any time after you create the table which will own the index. Also, you can create the index before or after you load data into the table. Generally, it is best to create the index before loading data into any table in a table space. If the table does not contain data, the system creates a description of the index. When you load the table, the system creates the index entries.

Index spaces

With DB2 and !DB/WORKBENCH, each index occupies its own index space. The index space consists of 1–64 VSAM datasets. When you create an index, the system automatically defines the index space in the same database as the table. !DB/WORKBENCH supports these VSAM data sets:

- A VSAM dataset that DB2 creates and manages
- A VSAM dataset that you have previously created and that you manage
- A VSAM dataset that you create when you create the index space

For information about which type of index space and the type of management of the VSAM dataset is best for your needs, see the *I O DB2 Administration Guide*.

Authorizations required to create an index space

To create an index, you must have one of these levels of authority:

- SYSADM authority
- DBADM authority for the database that contains the table to be indexed
- INDEX privilege for the table to be indexed
- Implicit authority because you created the table
- DBCTRL authority for the database (required only if the index name contains a qualifier that is not the same as the authorization ID)

General prerequisites for creating an index

You must have a VSAM dataset previously defined, or you must define one during these procedures. You also must meet one of these general prerequisites for creating an index:

- An existing database (!DB/WORKBENCH displays the default database DSNDB04 unless you have previously specified a database; however, you can type over the default value if you want.)
- An existing table space
- A table space which you create during the process of creating the index
- An existing table (with or without data)

Prerequisites common to unique, clustering, or partitioned indexes

Whether you use or modify the values from an existing index or from the defaults with !DB/WORKBENCH or whether you enter all new values, you must provide or confirm these values:

- Unique setting (for unique indexes only)
- Buffer pool name
- CLOSE setting
- Subpage setting
- Clustering index setting (for clustering or partitioned indexes)
- DEFER setting (If you want to defer the building of an index, you must provide the Y value for DEFER by using the O or the I options available from the Create Menu.)

Prerequisites for the VSAM dataset for the index space

When you create the index space, you also must set up values related to the VSAM dataset. You enter or confirm these values:

- ERASE setting
- Priqty value
- Secqty value
- PCTfree value
- FREEPAGE value
- Gbpcache value
- VSAM password
- Storage group name

Specific Tasks for Unique or Clustering Indexes

Overview

The unit contains a general overview of tasks required for creating unique or clustering indexes.

General overview of creating a unique index

With !DB/WORKBENCH you perform only one task to generate statements to create a unique index. When you create the index, you enter a value of U in the Unique field on the panel listing values for the new index.

General overview of creating a clustering index

With !DB/WORKBENCH, you perform only one task to generate statements to create a clustering index. When you create the index, you enter a value of Y in the Clustering Index? field on the panel listing values for the new index. (Although you can have multiple indexes on a table, you can only have one clustering index for each table.)

Specific Tasks for Partitioned Index and Index Space

Overview

The unit contains background information, a general overview of tasks, and prerequisites required for creating a partitioned index.

Partitioned table spaces and the effect on indexes

When you create a partitioned table space, you must create one index for the table in the table space. This index must be a partitioned and clustering index.

Because of this DB2 requirement, !DB/WORKBENCH automatically displays the Create Index panel after you create a table in a partitioned table space. Although you have more than one index on a table, you must create a clustering index as your first index on the table. (!DB/WORKBENCH, therefore, will not allow a change to the Y value in the Clustering Index? field.)

When you create a clustering index on a table in a partitioned table space, the index is also partitioned and is divided into multiple index spaces. You can assign each partitioned index to a different storage group.

General overview of creating a partitioned index

With !DB/WORKBENCH, you perform these tasks to generate statements to create a partitioned index and index space. When you set up the method you will use to create and manage the VSAM dataset, you perform one of two tasks:

- You specify that all partitions have the same value. (No other action is needed.)
 - You specify that all partitions have different values, and you complete panels specifying the values for each partitioned index space (by specifying values such as the ERASE setting and the Priqty value) and values for each partitioned index (by specifying the Limitkey values for each partition.)
-

Prerequisite specific to a partitioned index

If you create a partitioned index and want one or more partitions to have different values, you must specify the highest (or lowest if the column is used in descending order) value for each partition.

Beginning to Create the Index

Overview

The unit contains the instructions applicable to all options available when creating an index. It also contains preliminary instructions to follow if you accessed the Create Index panel from the create menu.

Prerequisite if you plan to manage the VSAM datasets

If you plan to manage the VSAM datasets, you must use a name for the index that is only 8 characters.

Preliminary steps to perform if you accessed the panel from the Create Menu

If you are creating an index for an existing table and if you accessed the Create Index panel from the Create Menu, follow these steps to enter the values needed to identify the database and the table and then follow the instructions on the next page to begin to create a new index.

Step	Action
1	On the Create Index panel, perform the appropriate action: <ul style="list-style-type: none"> ● Type the name of the database in the Database field. ● Type S in the Select field for the database you want to associate with the new index.
2	Press Enter.
3	Perform the appropriate action: <ul style="list-style-type: none"> ● Type the name of the creator of the table in the Table Creator field and the table name in the Table Name field. ● Type S in the Select field for the creator.table name you want to associate with the new index.
4	Press Enter. Results: The system displays a new panel with the Create Index field displayed.

Procedure to follow to begin to create an index

If you have just created a table or if you have performed preliminary steps on the previous page or you accessed the Create Index panel from the Create Menu, follow these steps.

Step	Action
1	If you want to change the Creator ID, type over the existing name with any Creator ID in the DB2 catalog.
2	Type the name of the new index in the Create Index field.
3	Press Enter. Results: The system displays a message directing you to select an index as a model and moves the cursor to the Select field for the first index on the list.

Determining what to do next

The chart lists the options !DB/WORKBENCH provides from the Create Index panel to create an index.

Characteristics of the New Table Space	Unit to Use
A new index that uses the same values as an existing index	“Using Values from an Existing Index” on page 294
A new index that uses default values (either from the system or from the default options you set up)	“Using the Default Values for the New Index” on page 296

Using Values from an Existing Index

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Index” on page 292 before you proceed with this unit, which covers one of the options available for creating an index.

This unit covers the information needed if you want to use the values from an existing index as the source of values for a new index. It describes the prerequisites and the procedures for this option, and gives instructions on what to do after you create the index.

Prerequisites for using values from an existing table space

To perform the procedures in this unit, you need the name of the existing table space you want to use as a model for the new index.

Reminder about unique, clustering, and partitioned indexes

If you are planning to create a unique, clustering, or partitioned index, be sure to confirm that the values in these fields are correct at this time.

- Unique field
- Clustering Index? field

Reminder about the DEFER value

If you are planning to defer the building of the index, be sure to confirm that a Y exists in the DEFER field.

Procedure for specifying an existing index

To specify an existing index, follow these steps.

Step	Action
1	On the Create Index panel, type S in the Select field for the index whose values you want to use for the new index.
2	Press Enter. Result: The system displays a panel showing the completed values for such fields as the Unique, buffer pool, CLOSE, Subpages and Clustering Index? fields, using the values in the model index you selected.
3	If needed, modify the values.
4	Press Enter. Results: The system displays a panel listing the columns on the table so that you can select the table columns to be indexed.

Determining what to do next

Follow the instructions in the unit “Selecting Table Columns to Be Indexed” on page 298.

Using the Default Values for the New Index

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Index” on page 292 before you proceed with this unit, which contains one of the three options available for creating an index.

This unit covers the information needed if you want to use default values from !DB/WORKBENCH or the default values you established using the default option (O on the Create Menu) for the new index you are creating. The unit describes the prerequisites and the procedures for this option, and gives instructions on what to do after you create the index.

Reminder about default values

If you need a reminder about how the system uses default values in this process, see “Introduction to Creating Objects” on page 209.

Reminder about unique, clustering, and partitioned indexes

If you are planning to create a unique, clustering, or partitioned index, be sure to confirm the values in these fields are correct at this time.

- Unique field
- Clustering Index field

Reminder about the DEFER value

If you are planning to defer the building of the index, be sure to confirm that a Y exists in the DEFER field.

Procedure for using default values for the new index

To use default values for the new index, follow these steps.

Step	Action
1	On the Create Index panel, type the name of the new index in the Create Index field.
2	Press Enter. Result: The system displays a list of indexes available as a model.
3	Press Enter again. Result: The system displays a panel showing the completed values for such fields as the Unique, buffer pool, CLOSE Subpages, Clustering Index? and DEFER fields using the default values.
4	If needed, modify values.
5	Press Enter. Result: The system displays a panel listing the columns on the table so that you select the columns to be indexed.

Determining what to do next

Follow the instructions in the unit “Selecting Table Columns to Be Indexed” on page 298.

Selecting Table Columns to Be Indexed

Overview

This unit contains background and the procedures for selecting table columns to be indexed.

Principles to follow when selecting table columns to be indexed

Follow these broad principles in selecting table columns.

- Specify no more than sixteen columns in an index.
 - Select columns that, when added, result in a total key length that is less than 40 bytes for a partitioned index or 250 bytes for other indexes. (The system displays your current total in the key length field on the line with the first column.)
 - Specify the columns in the order you prefer. (The order of the index columns can be a different order from the columns themselves.)
 - Specify the columns to be in either ascending or descending order.
-

Reminders about methods to select the table columns

You can use these methods to display and select the table columns.

- Use the Exclude (X) select in the A/D field to remove a column from the display or to remove a column you previously selected.
- Use the RESET command to redisplay all columns.

Interpreting data provided on the panel

The panel provides data about each table column to help you identify the column you want. In addition to the name and the data type of the column, the panel includes additional data about the table columns. The chart describes the data and the effect on the total key length if you select the column to be indexed.

Heading on the Panel	Contents	Effect on Key Length (if any)
Length	Length of the column	If you select a column with the VARCHAR data type, the system adds the full length of the key to the total for key length.
NL	Nullable column (NL indicates the column can be null.)	If you select a nullable column, the system adds 1 to the total for key length.
UP	Updateable column (NU indicates the column is not updateable.)	(Not applicable)

Procedure for selecting table columns to be indexed

Follow these steps when selecting table columns to be indexed.

Step	Action
1	Type this information for each table column you want to be indexed. A/D Whether the indexed field is in ascending (A) or descending (D) order Order The relative order of the index column
2	Press Enter. Result: The system adds the length of the column (and 1 byte if you have selected a nullable column) to the total key length.
3	When you have selected the last table column, press Enter. Result: The system adds the length of the last column to the total key length.
4	If the total key length is less than the maximum allowed, press Enter. Result: The system displays a panel showing the alternatives for creating and managing the VSAM dataset.

Determining what to do next

Review the chart to determine what to do next.

IF you created...	AND you want...	THEN use...
a unique or clustering index	DB2 to create and manage your VSAM dataset	the instruction in the unit “Completing a Non-Partitioned, DB2-Managed Dataset” on page 302
a partitioned index	DB2 to create and manage your VSAM dataset	the chapter “Completing Requirements for VSAM Datasets” on page 315
a unique, clustering, partitioned index	to manage the VSAM dataset which you previously created	the chapter “Completing Requirements for VSAM Datasets” on page 315
	to create a VSAM dataset at this time and to manage the dataset yourself	the chapter “Completing Requirements for VSAM Datasets” on page 315

Completing a Non-Partitioned, DB2-Managed Dataset

Overview

This unit contains instructions only if you have created a unique or clustering index and if you want DB2 to create and manage the VSAM dataset for you. This unit contains a reminder and the procedure to follow.

Reminder about the Space Estimator

You can use the !DB/WORKBENCH Space Estimator to estimate the value for Priqty. For an overview of the procedure to use the estimator, see the unit “Methods Available to Estimate Requirements for Space” on page 222 in the chapter “Introduction to Creating Objects” on page 209. For more detailed information, see “Using the DASD Estimators” on page 111.

Procedure to prepare a DB2-created and managed VSAM dataset

Follow these steps to set up a DB2-created and managed VSAM dataset.

Step	Action
1	On the panel listing alternatives for index space, type 1 in the Choose One field.
2	Press Enter. Result: The system displays a panel listing the values to be used for the VSAM dataset for the new table space.
3	If necessary, modify the values.
4	If the storage group name is not displayed, perform the appropriate action: <ul style="list-style-type: none"> ● Type the name of the storage group in the Storage Group field. ● Type S in the select field for the storage group you want to use.
5	Press Enter. Result: The system displays the Create Menu, or other options.

Determining what to do next

You have two broad choices from the Create Menu:

- If you want to create other objects, type the letter of the option in the Option field and press Enter.
- If you want to perform a task, such as to process the statements or save the statements from memory to a member in the user PDS, type **END** on the command line and press Enter, or press the appropriate function key to display the Process Member menu for the Create Menu.

If you need a reminder about using the Process Member menu from the Create Menu, see the chapter “Creating Objects” on page 207.

Introduction

This chapter explains prerequisites, options, and procedures for creating a storage group. The chapter “Introduction to Creating Objects” on page 209 is a prerequisite for this chapter.

Chapter contents

Overview of Chapter	306
Access to the Create Storage Group Panel	307
Prerequisites for Creating a Storage Group	308
Beginning to Create the Storage Group	309
Using Values and VOLSERS from an Existing Storage Group	310
Modifying Values from an Existing Storage Group	312

Overview of Chapter

Prerequisite for using this chapter

Before following the instructions in this chapter, review the chapter “Introduction to Creating Objects” on page 209

Organization of information in the chapter

The chapter provides background about accessing the Create Storage Group panel, prerequisites for creating a storage group (regardless of the method you use from the Create Storage Group panel), instructions for beginning the process of creating a storage group, and instructions for the two creation methods possible from the panel.

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Access	Prerequisites	Beginning	Model	Modified values
	p.307	p.238	p.309	p.310	p.312
Review methods of access (and the method to exit)	√				
Create a new storage group using an existing storage group as a model (without changing the values or the VOLSERS)		√	√	√	
Create a new storage group using an existing storage group as a model and modify the values or the VOLSERS		√	√		√

Access to the Create Storage Group Panel

Overview

The unit lists methods that result in display of the Create Storage Group panel and explains how to exit the panel if you do not want to create a new storage group at this time.

Methods to access the Create Storage Group panel

The system displays the Create Storage Group in these cases.

- You accessed the Create Storage Group panel from the Create Menu.
- You are creating another object that requires the entry of the name of a storage group and you enter a name of an object that is not defined in the DB2 catalog or is not currently stored in memory that is active during a session when you use the Create Menu.

For a reminder about instructions for the methods to access the Create Menu (including the fast path), see “Introduction to Creating Objects” on page 209.

If you do not want to create a storage group at this time

If you do not want to create a storage group at this time, follow these steps.

Step	Action
1	On the Create Storage Group panel, blank out the new name in the Create STOGROUP field.
2	Type END on the command line and press Enter, or press the appropriate function key.

Prerequisites for Creating a Storage Group

Overview

The unit contains the prerequisites required to create a storage group.

Terminology

With DB2 and !DB/WORKBENCH, a storage group is a set of DASD volumes. DB2 uses these volumes to allocate storage for table spaces and for indexes, to define VSAM datasets, and to manage the datasets (by extending and deleting the sets as needed).

Authorizations required to create a storage group

To create a storage group, you must have one of these levels of authority:

- SYSADM authority
 - CREATESG privilege
-

General prerequisites for creating a storage group

You must meet one of these general prerequisites for creating a storage group:

- An existing storage group to use as a model and a source for the list of previously defined VOLSERS
- VOLSERS that have previously been defined

Beginning to Create the Storage Group

Overview

The unit contains the instructions applicable to the two options available when creating a storage group.

Procedure to follow to begin to create the storage group

Follow these steps to begin to create the storage group.

Step	Action
1	On the Create Storage Group panel, type the name of the new storage group in the Create STOGROUP field.
2	Press Enter. Result: The system displays a message directing you to select a storage group to use as model and moves the cursor to the Select field for the first storage group in the list.

Determining what to do next

The chart lists the two options !DB/WORKBENCH provides from the Create Storage Group panel to create a storage group.

Characteristics of the New Storage Group	Unit to Use
A new storage group that uses an existing storage group as a model (without changing the values or VOLSERS)	“Using Values and VOLSERS from an Existing Storage Group” on page 310
A new storage group that uses an existing storage group as a model, but that modifies values or the list of VOLSERS	“Modifying Values from an Existing Storage Group” on page 312

Using Values and VOLSERS from an Existing Storage Group

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Storage Group” on page 309 before you proceed with this unit, which covers one of the two options available for creating a storage group.

This unit covers the information needed if you want to use the values from an existing storage group as the source of values for a new storage group (and do not plan to change any values from those in the model). It describes the prerequisites and the procedures for this option, and gives instructions on what to do after you create the storage group.

Prerequisites for using values from an existing storage group

To perform the procedures in this unit, you need to know the name of the existing storage group you want to use as a model (as the source of values for the VSAM catalog and the VOLSERS list) for the new storage group.

Procedure for using values from an existing storage group

To use values from an existing storage group, follow these steps.

Step	Action
1	On the Create Storage Group panel, type S in the Select field for the storage group whose values you want to use for the new storage group.
2	If you want, type a new name in the VSAM Catalog field.
3	Press Enter. Result: The system records the values from the model storage group in the VSAM Catalog field and displays the VOLSERS associated with the storage group.
4	Type N in the Edit Vollist? field.
5	Press Enter. Result: The system stores the values for the storage group, the VSAM Catalog, and the VOLSERS and redisplays the Create Menu.

Determining what to do next

You have two broad choices from the Create Menu:

- If you want to create other objects, type the letter of the option in the Option field and press Enter.
- If you want to perform a task, such as to process the statements or save the statements from memory to a member in the user PDS, type **END** on the command line and press Enter (or press the appropriate function key) to display the Process Member menu for the Create Menu.

If you need a reminder about using the Process Member menu from the Create Menu, see the chapter “Introduction to Creating Objects” on page 209.

Modifying Values from an Existing Storage Group

Overview

Be sure you have followed the instructions in the unit “Beginning to Create the Storage Group” on page 309 before you proceed with this unit, which contains one of the two options available for creating a storage group .

This unit covers the information needed if you want to modify values from an existing storage group (including if you want to modify a VOLSERS list) for the new storage group you are creating. The unit describes the prerequisites and the procedure for this option, and gives instructions on what to do after you create the storage group.

Background about editing the VOLSERS list

When you edit the VOLSER list when creating the storage group, you use the ISPF editor. Follow these specific rules for the VOLSER list:

- A VOLSER name must consist of 6 characters.
- The maximum number of VOLSERS is 133.

Prerequisites for modifying values from an existing storage group

To perform the procedure in this unit, you need to know the name of the existing storage group whose values you want to modify.

Beginning the procedure by specifying an existing storage group

To specify an existing storage group, follow these steps.

Step	Action
1	On the Create Database panel, type S in the Select field for the storage group you want to use as a model.
2	Press Enter. Result: The system displays the VOLSERS associated with the storage group you are using as a model.

Completing the procedure by editing the VOLSERS list

To modify values and edit the VOLSERS list, follow these steps.

Step	Action
1	If you want, type an new name in the VSAM Catalog field.
2	Type Y in the Edit Vollist? field
3	Press Enter. Result: The system invokes the ISPF editor and displays the VOLSERS from the model storage group for you to edit.
4	Edit the VOLSER list.
5	Type END on the command line and press Enter, or press the appropriate function key. Result: The system redisplayes the Create Storage Group panel and displays the VOLSERS.
6	Type N in the Edit Vollist? field
7	Press Enter. Result: The system stores in memory the CREATE statements using the values for the storage group (and its VOLSERS) and redisplayes the Create Menu.

Determining what to do next

You have two broad choices from the Create Menu:

- If you want to create other objects, type the letter of the option in the Option field and press Enter.
- If you want to perform a task, such as to process the statements or save the statements from memory to a member in the user PDS, type **END** on the command line and press Enter (or press the appropriate function key) to display the Process Member menu for the Create Menu.

If you need a reminder about using the Process Member menu from the Create Menu, see the chapter “Introduction to Creating Objects” on page 209.

Modifying Values from an Existing Storage Group

Chapter 14. Completing Requirements for VSAM Datasets

Introduction

This chapter explains prerequisites, options, and procedures for completing requirements for a VSAM dataset when creating a table space or index space. Completing the procedures in either the chapter “Creating a Table Space” on page 247 or the chapter “Creating an Index and Index Space” on page 283 is a prerequisite for this chapter.

Chapter contents

Overview of Chapter	316
Partitions and Setting Up the VSAM Dataset	320
Setting Up a DB2-Created and Managed Dataset	322
Setting Up an Existing Dataset That You Manage	325
Defining and Setting Up a VSAM Dataset	328
Background for a Partitioned Index Space	333
Completing Different Values for Partitions	335

Overview of Chapter

Prerequisite for using this chapter

Before following the instructions in this chapter, review the chapter “Introduction to Creating Objects” on page 209. You must also complete the procedures in either of these chapters:

- “Creating a Table Space” on page 247
- “Creating an Index and Index Space” on page 283

Background on table spaces and index spaces

The menus and panels you use when setting up a VSAM dataset for a table space or an index space are the same. The only exception is that the system displays an additional panel for each partition if you create a partitioned index space and if you want one or more partitions to have different values. The additional panel allows you to input the value for the highest index key value for each partition (or lowest value if you specified descending order). When you have learned to set up a VSAM dataset for a new table space, you have also learned how to use most of the menus and panels you will use when you set up a new index space.

Organization of information in the chapter

The chapter provides prerequisites and procedures for completing requirements for VSAM datasets required for new table spaces or index spaces if you are creating one of these types of datasets using options available with the Create Menu.

- A DB2-created and managed VSAM dataset for a partitioned table space or index space
- An existing VSAM dataset that you plan to manage and you want to use for a table space or index space that you have created
- A VSAM dataset that you plan to define by using options available with !DB/WORKBENCH and that you plan to manage and you want to use for a table space or index space you have created

Organization of information and your needs

This unit includes three charts showing how to select units of instruction for the three options for managing the VSAM dataset required for table spaces and index spaces available from !DB/WORKBENCH.

Units to use if you want DB2 to create and manage the dataset

If you want DB2 to create and manage the VSAM dataset for the table space or index space, review the chart to select information appropriate for the task you want to perform. (Instructions for simple and segmented table spaces and for indexes that are not partitioned are covered in the chapters “Creating a Table Space” on page 247 and “Creating an Index and Index Space” on page 283.)

	Partition	DB2 Created	Existing Dataset	Defined Dataset	Partitioned Index	Different Partitions
	p.320	p.322	p.325	p.328	p.333	p.335
Complete a partitioned table space with all partitions having the same values	√	√				
Complete a partitioned index space with all partitions having the same values	√	√				
Complete a partitioned table space with one or more partitions having different values	√	√				√
Complete a partitioned index space with one or more partitions having different values	√	√			√	√

Units to use if you want to set up and manage an existing dataset

If you want to manage the VSAM dataset and you have an existing VSAM dataset for the table space or index space, review the chart to select information appropriate for the task you want to perform.

	Partition	DB2 Created	Existing Dataset	Defined Dataset	Partitioned Index	Different Partitions
	p.320	p.322	p.325	p.328	p.333	p.335
Complete a simple or segmented table space			√			
Complete a unique or clustering index			√			
Complete a partitioned table space with all partitions having the same values	√		√			
Complete a partitioned index space with all partitions having the same values	√		√			
Complete a partitioned table space with one or more partitions having different values	√		√			√
Complete a partitioned index space with one or more partitions having different values	√		√		√	√

Units to use to define and set up a VSAM dataset

If you want to manage the VSAM dataset for the table space or index space and if you want to define the dataset during this session, review the chart to select information appropriate for the task you want to perform.

	Partition	DB2 Created	Existing Dataset	Defined Dataset	Partitioned Index	Different Partitions
	p.320	p.322	p.325	p.328	p.333	p.335
Complete a simple or segmented table space				√		
Complete a unique or clustering index				√		
Complete a partitioned table space with all partitions having the same values	√			√		
Complete a partitioned index space with all partitions having the same values	√			√		
Complete a partitioned table space with one or more partitions having different values	√			√		√
Complete a partitioned index space with one or more partitions having different values	√			√	√	√

Partitions and Setting Up the VSAM Dataset

Overview

This unit is applicable only if you created a table space or index space. The unit contains an overview of the panel that the system displays when you assign values to the partition and provides options for using this panel.

Overview of the panel for assigning values to partitions

When you select the option to have DB2 create and manage your VSAM datasets for a partitioned table space or index space, the system displays the panel and highlights the message indicating the number of the partition and how many partitions you have defined.

```

NEW TBLSPC OPTIONS -      - DB/WORKBENCH --- DB2=DB2 ----- STOGROUP 1 OF 58
Cmd ==>>                                     Scroll ==>> PAGE

                C R E A T E   T A B L E S P A C E

Fill in the values and press ENTER.  Partition 1 of 2
Erase           ==>> N                Y/N  Erase disk when delete dataset?
Priqty          ==>> 12               Primary Allocation (in KB). ? to Est.
Secqty          ==>> 4                Secondary Allocation (in KB).
Pctfree        ==>> 5                (0 - 99) Pct freespace per page.
Freepage       ==>> 0                (0 - 255) Pages between free pages.
Gbpcache       ==>>                 A/blank Group buffer pool caching.
Password       ==>>                 VSAM password
Storage Group  ==>> SYSDEFLT         (or select from list below)
                                           Type REV to go back 1 panel.
-----

```

STOGROUP	STOGROUP	STOGROUP	STOGROUP	STOGROUP	STOGROUP
- CILESG	- DADDSTG	- DB2BASE	- DDAYSTG	- DHDBTSG1	- DHDBTSG3
- DHG31E	- DHSG31	- D097STG	- DSN8GJP0	- DSN8G31U	- DSN8G310
- D097STG	- D1219STG	- GULSTO	- GV000	- GV001	- HOUGROUP
- JFSTGRP	- KTCSTG	- KTETG001	- KTETG002	- KTETG003	- KTS01SGT
- KTS17SGT	- LOANDF	- LRISHDSG	- LRISTG1	- LRITTOOLS	- MMMSTO
- MMMSTO2	- PATERNSG	- PPCDBS	- QASTG	- REGEXP	- REHSTO
- RUTHSTO	- SGMARIA	- SGSEL01	- SGTSBN03	- SGTSLA66	- SMUTBED
- SMUTSTSG	- STGCILE	- STITHSTG	- STOGROUP	- SYSDEFLT	- TDDB18SG
- TDDB18S2	- TSGM35SG	- TSL238	- TSL238B	- TSL238F	- TSUKSTG
- UK87DB31	- UK87PLEX	- V240SG	- ZETTI	*** END ***	

Options for responding to the panel

Review the chart for options for responding to the panel and the system response.

Condition	Action of the System
You plan for all partitions to have the same values	The system applies the values you accept or modify on the panel to each partition in the VSAM dataset (not just the first partition).
You plan for one or more partitions to have different values	The system applies the values you accept or modify only to the first partition. (The system then redisplay the panel with the message indicating the number of the next partition.)

If you make an error

When setting values for partitions (unlike other panels used when creating objects), you cannot use the REV command to return to a prior panel. If, however, you make an error, you can edit the CREATE statements.

Setting Up a DB2-Created and Managed Dataset

Overview

This unit contains background prerequisites, the procedure to follow if you want to set up a DB2-created and managed VSAM dataset for a partitioned table space or index space and instructions about what to do after you complete the procedure. Be sure you have reviewed the unit “Partitions and Setting Up the VSAM Dataset” on page 320.

Reminder about the Space Estimator

You can use the !DB/WORKBENCH Space Estimator to estimate the value for Priqty. For an overview of the procedure to use the estimator, see the unit “Methods Available to Estimate Requirements for Space” on page 222. For more detailed information, see “Using the DASD Estimators” on page 111.

Reminder about instructions for non-partitioned datasets

If you are creating a simple or segmented table space or a unique or clustering index that is not partitioned, instructions for preparing a DB2-managed dataset are included at the end of each of these chapters:

- “Creating a Table Space” on page 247
- “Creating an Index and Index Space” on page 283

Prerequisites for the VSAM dataset

When you create the table space or index space, you also must set up values related to the VSAM dataset. You enter or confirm these values:

- ERASE setting
- Priqty value
- Secqty value
- PCTFREE value
- FREE page value
- VSAM password
- Storage group name

Procedure to set up a DB2-created and managed VSAM dataset

Follow these steps to set up a DB2-created and managed VSAM dataset for partitioned table spaces or index spaces.

Step	Action
1	On the panel listing alternatives for table spaces or index spaces, type 1 in the Choose One field.
2	Press Enter. Result: The system displays a panel listing the values to be used for the VSAM dataset for the new table space or index space and a message indicating that these values are for the first partition.
3	If necessary, modify the values.
4	If the storage group name is not displayed, perform the appropriate action: <ul style="list-style-type: none"> ● Type the name of the storage group in the Storage Group field. ● Type S in the select field for the storage group you want to use.
5	Press Enter. Result: The system displays a panel with the prompt asking whether the characteristics of the partitions are the same.
6	Perform the appropriate action: <ul style="list-style-type: none"> ● If the partitions are to have the same values, press Enter. ● If the partitions are to have different values, type Y over the N in the field and press Enter. Result: The result varies with your previous action: <ul style="list-style-type: none"> ● If you entered N, the system applies the values you used for the first partition to all partitions and then displays a menu or panel for optional tasks (such as creating dependent objects). ● If you entered Y, the system displays the panel listing the values to be used for the second partition.

Determining what to do next

If you want all partitions to have the same values, no other action is necessary. You can create another object or process the statement you have generated. (If you have created a partitioned table space, you should also create a table and a partitioned index.) If you need a reminder about how to process statements, see the chapter “Introduction to Creating Objects” on page 209.

If you want one or more partitions to have different values, review the chart for the instructions to use depending on the object you have been creating.

Object You Are Creating	Source to Use
Table Space	“Completing Different Values for Partitions” on page 335
Index Space	“Background for a Partitioned Index Space” on page 333

Setting Up an Existing Dataset That You Manage

Overview

The unit contains background, prerequisites, and the procedure to follow if you want to set up an existing VSAM dataset for a table space or index space that you will manage. It includes instructions for partitioned and non-partitioned datasets. (If you are creating a partitioned table space or index space, be sure you have reviewed the unit “Partitions and Setting Up the VSAM Dataset” on page 320.)

Prerequisites if you want to manage an existing VSAM dataset

If you want to manage an existing VSAM dataset, you enter or confirm these values:

- VSAM Catalog name
- VSAM password
- PCTFREE value
- FREE page value

Procedure to set up an existing dataset that you manage

Follow these steps if you want to manage the dataset yourself and you previously created it.

Step	Action
1	On the panel listing alternatives for tables space or index spaces, type 2 in the Choose one field.
2	Type 1 in the Select field.
3	Press Enter. Result: The system displays a panel listing the values to be used for the VSAM dataset for the new table space or index space. (If you are creating a partitioned dataset, the system displays a message indicating that these values are for the first partition.)
4	If necessary, modify values on the panel.
5	Press Enter. Result: The result varies with your previous action: <ul style="list-style-type: none"> ● If you are not creating a partitioned dataset, the system completes the set up of the dataset and displays a menu or panel for optional tasks (such as creating dependent objects). ● If you are creating a partitioned dataset, the system displays a panel with a prompt asking whether the characteristics of the partition are the same.
6	If you are creating a partitioned dataset, perform the appropriate action: <ul style="list-style-type: none"> ● If the partitions are to have the same values, press Enter. ● If the partitions are to have different values, type Y over the N in the field and press Enter. Result: The result varies with your previous action: <ul style="list-style-type: none"> ● If you entered N, the system applies the values you used for the first partition to all partitions and then displays a menu or panel for optional tasks (such as creating dependent objects). ● If you entered Y, the system displays the panel listing the values to be used for the second partition.

Determining what to do next

If you are not creating a partitioned dataset or if you want all partitions to have the same values, no other action is necessary. You can create another object or process the statement you have generated. (If you have created a partitioned table space, you should also create a table and a partitioned index.) If you need a reminder about how to process statements, see the chapter “Introduction to Creating Objects” on page 209.

If you want one or more partitions to have different values, review the chart for the instructions to use depending on the object you have been creating.

Object You Are Creating	Source to Use
Table Space	“Completing Different Values for Partitions” on page 335
Index Space	“Background for a Partitioned Index Space” on page 333

Defining and Setting Up a VSAM Dataset

Overview

The unit contains background about processing and editing AMS DEFINE statements and about the storage group name, prerequisites, the procedures to follow if you want to define, during this procedure, a VSAM dataset, and you will manage the dataset for a new table space or index space. It includes instructions for partitioned and non-partitioned datasets. (If you are creating a partitioned table space or index space, be sure you have reviewed the unit “Partitions and Setting Up the VSAM Dataset” on page 320.)

Background about options available for DEFINE statements

The chart lists the options available for the AMS DEFINE statements used during the process of creating the table space or index space. (The options available are different from the options available when you process statements generated by the AMS command because the statements generated from the Create Menu are initially stored in memory and not in a member of the user PDS.)

Options on the Menu	System Response
DEFINE Now	Invokes AMS and defines the VSAM datasets at this time
EDIT	Invokes the ISPF editor and displays the DEFINE statements (When you complete editing, the system redisplay the DB2 Dataset Allocation panel.)
CANCEL	Cancel the VSAM dataset definitions that are stored in memory
PRINT	Prints the AMS DEFINE statements
END	Saves the changes in a member of the user PDS (The member name consists of the USERID followed by the letter A. The next time you create a table space or index space, the system overwrites the contents of the member.)

If you choose to edit the AMS DEFINE statements

If you choose to edit the AMS DEFINE statements, you should change only these values. The chart lists the values and an example from the statements.

Value You Can Change	Example from the Statement
VOLSER list	VOLUMES (VL0196)
Allocation amounts	Records (10 10)
ERASE/NOERASE	NOERASE

Prerequisites if you want to define and manage the VSAM dataset

If you want to define and manage the VSAM dataset for a new table space or index space, you enter or confirm these values:

- VSAM Catalog name
- VSAM password
- ERASE setting
- Priqty values
- Secqty value
- PCTFREE value
- FREE page value
- Storage group name

Reminder about the Space Estimator

When you create the table space or index space, you also must set up values related to the VSAM dataset. You enter or confirm these values:

- ERASE setting
- Priqty value
- Secqty value
- PCTFREE value
- FREE page value
- VSAM password
- Storage group name

Background about the storage group name

When defining the VSAM dataset, the storage group is used to name a collection of VOLSERS that are required when defining the dataset. You do not have to have DB2 authorization to use the storage group because the system does not use the storage group, but instead extracts the list of VOLSERS. You can use any storage group that is defined to DB2 to provide the list of VOLSERS.

You have two options for entry on the panel when defining the storage group name as the source of the VOLSERS. The chart lists your options and the system response when you press the Enter key.

Your Option	System Response
Type the name of a new storage group in the Storage Group field.	<p>Uses the VOLSERS listed for the storage group when generating the AMS DEFINE statement</p> <p>Automatically generates a CREATE STOGROUP statement (You do not need to use the Create Storage option on the Create Menu, and you can delete the statement if you want.)</p>
Type an S in the Select field for the storage group to be used as a model.	<p>Uses the VOLSERS listed for the storage group when generating the AMS DEFINE statement</p> <p>Automatically generates a CREATE STOGROUP statement (You do not need to use the Create Storage option on the Create Menu, and you can delete the statement if you want.)</p>

Procedure to define and set up the VSAM dataset

Follow these instructions if you want to define and set up a VSAM dataset that you plan to manage.

Step	Action
1	On the panel listing alternatives for table spaces or index spaces, type 2 in the Choose one field.
2	Type 2 in the Select field.
3	Press Enter. Result: The system displays a panel listing the values to be used for the VSAM dataset for the new table space or index space. (If you are creating a partitioned dataset, the system displays a message indicating that these values are for the first partition.)
4	If necessary, modify values on the panel.
5	Press Enter. Result: The result varies with your previous action: <ul style="list-style-type: none"> ● If you are not creating a partitioned dataset, the system displays the DB2 Data Set Allocation panel. ● If you are creating a partitioned dataset, the system displays a panel with a prompt asking whether the characteristics of the partitions are the same.

Determining what to do next

Review the chart to determine what to do next. (If you have created a partitioned table spaces, you should also create a table and a partitioned index.)

IF you are . . .	THEN . . .
Not creating a partitioned dataset	<p>On the DB2 Data Set Allocation Menu, type the letter of the option you want and press Enter.</p> <p>Result: The system processes the statements according to your choice and displays a menu or panel.</p>
Creating a partitioned dataset and want all partitions to have the same values	<p>On the panel asking whether the characteristics of the partitions are the same, press Enter.</p> <p>Result: The system displays the DB2 Data Set Allocation Menu.</p> <p>Type the letter of the option you want and press Enter.</p> <p>Result: The system processes the statements according to your choice and displays a menu or panel.</p>
Creating a partitioned dataset and want one or more partitions to have different values	<p>On the panel asking whether the characteristics of the partitions are the same, type N over the Y in the field, and press Enter.</p> <p>Results: The result varies with your action:</p> <ul style="list-style-type: none"> ● With a partitioned table space, the system displays the panel listing the values for the second partition. ● With a partitioned index space, the system displays a panel listing the columns that you have selected as keys for your index for the table. <p>Follow the instruction in the unit “Completing Different Values for Partitions” on page 335.</p>

Background for a Partitioned Index Space

Overview

If you are setting up a partitioned index space, review this unit before beginning the unit “Completing Different Values for Partitions” on page 335. This unit contains background needed if you are completing values for different partitions.

Order of panels and setting up a partitioned index space

After you specified values for the first partition (1 of n) and after you specified that you wanted one or more partitions to have different characteristics, the system displays a panel for you to use to specify the limit values for the key columns for the first partition. When you have specified those values, the system displays the panel for the values for the index spaces for the second partition.

Background if you are setting up a partitioned index space

If you are setting up a partitioned index space and want one or more partitions to have different values, you must specify the beginning and ending keys for each partition. (DB2 supplies the key values for the last partition; therefore, the system does not display a panel for the key range for the last partition.)

!DB/WORKBENCH provides a panel showing the key columns for the new index you created and indicating the number of the partition for which you are establishing the limit keys. You do not need to enter a value for each column, but you must enter a value in at least one column.

Type of Data	Action to Perform
Numeric (such as INT or FLOAT)	Do not enclose.
String (such as CHAR or VARCHAR)	Enclose in quotes. Type as lower case or upper case. (The system does not automatically capitalize letters.)

How to set up the index key values for each index space partition

If you are setting up a partitioned index space, the system requires you to provide a high value for each partition. (You must also set up the values for the VSAM dataset to be used for each partition.) Follow these steps for each partition for which you need to provide values.

Step	Action
1	On the panel listing the key columns you previously selected for the new index, type the highest value for the key for the system to use as the limit key for the partition.
2	Press Enter.
3	If you want, type the highest value for another key for the system to use as a limit key for the partition.
4	Press Enter.
5	When you have specified all the values for the partitions, press Enter again. Result: The result varies with the member of the partition: <ul style="list-style-type: none">● If you need to define the values for another partition, the system displays a panel with a message indicating the number of the partition and displaying the values required.● If you have defined all partitions, the system displays a menu or panel for optional tasks, such as creating dependent objects (if you are setting up an existing dataset that you will manage) or the DB2 Data Set Allocation Menu (if you are defining the VSAM dataset).

Completing Different Values for Partitions

Overview

The unit is applicable only if you want different values for one or more partitions for table spaces or index spaces. (If you are completing values for a partitioned index space, review the unit “Background for a Partitioned Index Space” on page 333 before beginning this unit.)

How to complete entry of different values for partitions

Follow these steps for each partition for which you need to provide values.

Step	Action
1	If you are setting up a partitioned index space, complete the limit key values for the index keys for the first partition and press Enter.
2	On the panel with a message indicating that the values are for the second partition, modify or complete the value if necessary.
3	Press Enter. Result: The result varies with your previous action: <ul style="list-style-type: none"> ● If you are setting up a partitioned table space, the system displays the panel indicating that values are for the third partition. ● If you are setting up a partitioned index space, the system displays the panel for the limit key values for the second partition.

Completing entry of values for the partition

When you have specified values for all partitions and limit keys for all partitions but the last (if you are setting up a partitioned index space), the system generates the statements and stores them in memory. If you are defining a VSAM dataset, the system also displays the DB2 Data Set Allocation Menu.

If the system displays the DB2 Data Set Allocation Menu

If the system displays the DB2 Data Set Allocation Menu, follow these steps to define the new VSAM dataset.

Step	Action
1	On the DB2 Data Set Allocation Menu, type the letter of the option you want.
2	Press Enter. Result: The system processes the statement according to your selection and displays a menu or panel for optional tasks (such as creating dependent objects).

If you need a reminder about the DB2 Data Set Allocation Menu, see “Defining and Setting Up a VSAM Dataset” on page 328.

Determining what to do next

You have two broad choices from the Create Menu:

- If you want to create other objects, type the letter of the option in the Option field and press Enter. (If you have created a partitioned table space, you should also create a table and a partitioned index.)
- If you want to perform a task, such as to process the statements or save the statements from memory to a member in the user PDS, type **END** on the command line and press Enter (or press the appropriate function key) to display the Process Member menu for the Create Menu.

If you need a reminder about using the Process Member menu from the Create Menu, see the chapter “Introduction to Creating Objects” on page 209.

Running Utilities

Chapter 15. Introduction to DB2 Utilities and Options

Introduction

This chapter provides background and instructions for using DB2 utilities and other options with JCL that are available with !DB/WORKBENCH. This chapter is a prerequisite for the other chapters in this section.

Chapter contents

Resources for Learning about Topics in This Section	340
Overview of the Chapter	342
Process for Using Utilities and Profiles	343
Accessing Object List Panels to Use with Utilities	344
Accessing RI Data to Use with Utilities	346
Introduction to the Utilities Menus	348
Introduction to the List Option	349
Introduction to Utility Profiles and !DB/WORKBENCH	350
Setting the Active Utility Profile	351
Accessing the Utilities	354
Completing Values for the Utility	356
Processing the Statements Generated	359
Accumulating Statements from Multiple Utilities	361

Resources for Learning about Topics in This Section

Overview

This unit lists resources available for learning about DB2 utilities and other options available with !DB/WORKBENCH. It includes information about the resources in this guide and in !DB/WORKBENCH.

Background about !DB/WORKBENCH and DB2 utilities and options

!DB/WORKBENCH provides easy access to DB2 utilities and other options either by using commands or from menus for utilities for different types of objects. !DB/WORKBENCH enhances the utilities by building a complete and executable job stream that you can submit from a Process Member menu. The statements that !DB/WORKBENCH builds include the JOB card, EXEC and DD statements, and necessary utility control statements. In addition, the system generates DD statements tailored according to your installation standards.

With !DB/WORKBENCH, you have substantial control over how the utilities run without extensive keying or experience with JCL. For example, you can perform these types of tasks:

- Generate a job stream that affects one object or all objects displayed on the object list panel
- Perform utilities on the entire table space or index space or on a specific partition
- Dynamically generate the utility job stream for an entire database even when users are constantly updating with new table spaces and indexes
- Use the default utility profile or create your own utility profiles for different tasks
- Run the utilities online or use batch processing

Resources in this section of the guide

This guide is intended to provide a broad picture of how to use the utilities available with !DB/WORKBENCH. The chart lists sources available within this section.

Information You Want	Resources Available
Introductions to accessing and using the utilities and on setting active utility profiles	This chapter
Introduction to each utility, including background and specific values required to use the utility	“DB2 Utilities and Options Available with !DB/WORKBENCH” on page 363
Introduction to fields and panels required to use multiple utilities	“Panels and Fields Used with Multiple Utilities” on page 419
Instructions for creating and modifying a utility profile	“Creating or Modifying Profiles for DB2 Utilities and Options” on page 449
Instructions for creating user variables for the utility profile	“Using User Variables in Utility Profiles” on page 463

Resources in the next section of this guide

This guide also contains information that you need when using utilities online and also in batch. This information common to both online and batch is available in “JCL for Utilities, Batch Jobs, and Extracts” on page 521.

The chapter provides this type of information:

- Overview of the skeletons and the information included in the skeleton
- Overview of how !DB/WORKBENCH uses the values you provide on panels (or in the utility profile) to substitute for values in the statements !DB/WORKBENCH generates

Overview of the Chapter

Organization of information in the chapter

The chapter provides background and instructions for using DB2 utilities and other options with !DB/WORKBENCH. It also contains information essential to the other chapters in this section and to using any utility described in the section.

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Accessing Objects	Accessing RI	Utilities Menus	List Option	Utility Profiles	Setting Active Profiles	Accessing Utilities	Completing Values	Processing Statements	Accumulating Statements
	p.344	p.346	p.348	p.349	p.350	p.351	p.354	p.356	p.359	p.361
Use utilities with object lists such as table spaces, indexes, and storage groups	√									
Use utilities specifically with RI (referential integrity) data		√								
Review the options available with utilities menus			√							
Review how to exclude objects from processing within the utility				√						
Review how utility profiles work and how to set the active profile					√	√				
Start a utility from the command line or a utility menu profiles							√			
Complete the panels required for a utility								√		
Execute the statements stored in the process member or use other options from the Process Member menu									√	
Accumulate statements generated by more than one utility into the process member										√
Use utilities for the first time	√	√	√	√	√	√	√	√	√	√

Process for Using Utilities and Profiles

Overview

This unit introduces the broad steps you follow when using a utility.

Background about the process

When you use a utility online, you follow these broad steps. (If you use the DIAG ABENDS utility, the utility takes effect without displaying the Process Member menu.)

Step	Action
1	Access the objects that you want to use with the utility.
2	Set (make active) the utility profile you want to use.
3	Access the utility.
4	Complete fields or panels that are required for the utilities you want.
5	Use the Process Member menu to submit, edit, delete, or print the utility statements you have generated.

Prerequisite for using a default utility profile

!DB/WORKBENCH provides a default utility profile that it uses if you do not set your own active profile. If you want to use a default utility profile, see “DB2 Utilities and Options Available with !DB/WORKBENCH” on page 363.

Accessing Object List Panels to Use with Utilities

Overview

This unit contains information on the availability of utilities on the Table Spaces, Indexes, and Storage Groups object list panels.

Availability of DB2 utilities and other options

The chart shows the availability of specific utilities on the object list panels.

	Table Spaces	Indexes	Storage Groups
CHECK DATA	√		
CHECK INDEX	√	√	
COPY	√		
DIAG ABENDS	√	√	
DIAG SYSUTIL	√		
DISASTER REC	√		
ICCAT	√		
MERGE COPY	√		
MODIFY	√		
QUIESCE	√		
REBIND	√	√	
REBIND PKGS	√	√	
REBIND PLANS	√	√	
RECOVER	√	√	
REORG	√	√	
REPORT	√		
RUNSTATS	√	√	
STOSPACE			√

Using utilities from the command line or a menu

With !DB/WORKBENCH, you can use utilities from either the command line or from a menu. For example, if you use the UTIL command or the Util (U) select on the Table Spaces panel, !DB/WORKBENCH displays the Table Spaces Utilities menu, listing utilities such as ICCAT, CHECK DATA, and RUNSTATS.

This chapter and access from the command line or menu

This chapter includes instructions for access from the command line or from a menu that you can display. In each case, the information about using the utilities from the command line is covered first.

Implications for macros and batch processing

If you want to use utilities in macros or in batch processing, you use the same commands and parameters for utilities available from the command line. (One exception exists: You cannot use the GEN parameter with macros or batch.)

Determining what to do next

Once you are on the object list panel for the objects you want, see “Setting the Active Utility Profile” on page 351. If you have not used options available for a utility menu previously, also see “Introduction to the Utilities Menus” on page 348, “Introduction to the List Option” on page 349, and “Introduction to Utility Profiles and !DB/WORKBENCH” on page 350.

Accessing RI Data to Use with Utilities

Overview

This unit provides methods for accessing RI data so that you can use utilities to manage referential data.

Background about table space sets

When you access RI data in !DB/WORKBENCH, it displays the data in table space sets (or in lists of table spaces that contain tables with referential relationship). The two object list panels with RI data listed as table space set provides similar data. The two panels, however, provide slightly different data:

- Using the TSSET command, you see the names of tables and you see the names of tables in the set.
 - Using the RI (F) select and the SET command, you see quantities of tables and indexes.
-

Accessing RI data beginning with the RI (F) select

Follow these steps.

Step	Action
1	On the Tables panel, type F in the Select field for a table with referential relationship.
2	Press Enter. Result: The system displays the Referential Relations panel which supports the SET command.
3	Type SET on the command line.
4	Press Enter. Result: The system displays the Referential TS Set panel which supports the UTIL command.

Accessing RI data beginning with the TSSET command

Follow these steps. (You can use filtering to display the object list you want.)

Step	Action
1	On the Tables panel, type TSSET on the command line.
2	Press Enter. Result: The system displays the Referential TS Set panel which supports the UTIL command.

Using utilities from the command line or a menu

With !DB/WORKBENCH, you can use utilities from either the command line or from a menu. For example, if you use the UTIL command on a Referential TS Set panel, !DB/WORKBENCH displays the Table Spaces Utilities menu, listing utilities such as ICCAT, CHECK DATA, and RUNSTATS.

This chapter and access from the command line or menu

This chapter includes instructions for access from the command line or from a menu that you can display. In each case, the information about using the utilities from the command line is covered first.

Implications for macros and batch processing

If you want to use utilities in macros or in batch processing, you use the same commands and parameters for utilities available from the command line. (One exception exists: You cannot use the GEN parameter with macros or batch.)

Determining what to do next

Once you are on the object list panel for the objects you want, see “Setting the Active Utility Profile” on page 351. If you have not used options available for a utility menu previously, also see “Introduction to the Utilities Menus” on page 348, “Introduction to the List Option” on page 349, and “Introduction to Utility Profiles and !DB/WORKBENCH” on page 350.

Introduction to the Utilities Menus

Overview

This unit covers the elements in the utilities menus that you can display from the object list panels by using the UTIL command or (on some panels) the Util (U) select.

Elements of the utilities menu

The utilities menu (whether for table spaces or other objects) is organized in the same way. The variation is in the number of utilities available for different objects.

```

----- TABLESPACE UTILITIES -----
Option ==> 1          SET PROFILE NAME ==> DUMMY 5

 0 LIST                List tablespaces to exclude some. 2
 1 CHECK INDEX         Check all Indexes in a Tablespace.
 2 CHECK DATA         Check Referential Relationships.
 3 COPY                Copy Data into a Sequential Data Set.
 4 MERGECOPY           Merge Multiple Copies into one Copy.
 5 MODIFY              Purge records from SYSIBM.SYSCOPY.
 6 QUIESCE             Stop Changes to the Tablespaces.
 7 RECOVER             Recover a Tablespace.
 8 REORG               Reorganize a Tablespace.
 9 REPORT              Report Recover Status on Tablespaces. 3
10 RUNSTATS           Get Tablespace and Index Statistics.
11 REBIND              Rebind Dependent Plans and Packages.
12 REBIND PLANS        Rebind Dependent Plans.
13 REBIND PKGS         Rebind Dependent Packages.
14 ICCAT               Catalog Image Copy Datasets.
15 PROFILE             Create/Modify Utility Profiles. 4
16 DIAG SYSUTIL        Generate DIAGNOSE SYSUTIL.
17 DIAG ABENDS         Generate DIAGNOSE ABEND statements. 3
18 DISASTER REC        Recover a Lost Tablespace
  
```

- 1** Field for entry of options for the utility
- 2** Option that lets you exclude table spaces from processing
- 3** Utilities available for the type of object
- 4** Option that lets you display a list of existing utility profiles for selection (You can also create, copy, display, or modify utility profiles from this list.)
- 5** Option that lets you set the active utility profile

Introduction to the List Option

Overview

This unit reviews the List option that is available from utilities menus and that allows you to exclude objects from processing.

Using the List option from the utilities menu for the object list panel

You can restrict the number of objects affected by the utility by using the List option from the utilities menu of the object list panel. Follow these steps.

Step	Action
1	On the utilities menu for the object list panel, type 0 in the Option field.
2	Press Enter. Result: The system displays a panel listing the objects to be affected by the utility. (If you have previously excluded an object, the system displays it with highlighted text.)
3	Type any character in the Select field for each object you want to exclude from processing.
4	Type END on the command line and press Enter, or press the appropriate function key. Result: The system excludes from processing those objects that you have excluded.

If you want to exclude partitioned table spaces

The list that the system displays shows only the first partition of any partitioned table space. If you select the first partition for exclusion, the system excludes the entire table space.

If you want to exclude a single partition, do not select it on the panel displayed from the List option. Instead, specify the number of the partition in the DSNUM/PART field shown on the Partition Specifications panel. (The system automatically displays the panel if you select an object with partitions and if the utility you are using supports partitions.)

Introduction to Utility Profiles and !DB/WORKBENCH

Overview

This unit provides introductory information about utility profiles and information about how !DB/WORKBENCH uses utility profiles.

!DB/WORKBENCH and the default profile

You can use utilities available with !DB/WORKBENCH without creating your own utility profile. !DB/WORKBENCH provides a default profile (named Dummy) that is used if you cannot make another profile active (or set the profile).

!DB/WORKBENCH saves the utility profiles in the system PDS. Only users who are authorized to update the system PDS can create and save a personal utility profile.

Utility profiles and your USERID

!DB/WORKBENCH associates the default utility profile or the profile you set as your active profile with your USERID. In this way, when you begin a session with !DB/WORKBENCH, the utility profile that you used last is active.

Contents of a utility profile

Utility profiles consist of a series of panels containing values you want to use. The profile must contain these values:

- Values specific to the utility or option you want to use (such as values required for the REORG utility)
 - Values used by the utility you want to use as well as by other utilities (such as the job options or the dataset information)
-

Active and inactive utility profiles

You can have as many utility profiles as you need for the way you want to run utilities. You can, however, have only one active (set) utility profile.

Setting the Active Utility Profile

Overview

This unit contains background and instructions for three methods for setting the active utility profile. (The unit “Introduction to the Utilities Menus” on page 348 is a prerequisite for this unit.)

Methods for setting the active profile

This unit contains instructions for setting the active profile:

- Directly from the command line
- From a utilities menu
- From a list that you can display from a utilities menu

For information about creating and modifying profiles, see “Creating or Modifying Profiles for DB2 Utilities and Options” on page 449.

Prerequisite for all procedures in this unit

To perform the procedures in this unit, you must begin your actions on an object list panel that supports the UTIL command or the Util (U) select.

Background about parameters for the the UTIL command

You can use the OPTS parameter followed by the name of the utility profile as parameters for the UTIL command. In this way, you can set the active utility profile without going through a utilities menu. For example, on the Table Spaces panel, you can set the active profile to your profile named *TEST* by using this command on the command line.

```
UTIL OPTS TEST
```

Setting the active utility profile from the command line

Follow these steps to set the active utility profile directly from the command line (without displaying a utilities menu).

Step	Action
1	On the command line for an object list panel, type UTIL OPTS <i>profile name</i> on the command line.
2	Press Enter. Result: The system sets the utility profile. (The system does not display a message.)

Setting the active utility profile on a utilities menu

Follow these steps to specify the active utility profile on a utilities menu.

Step	Action
1	On the utilities menu for an object list panel, type the name of the utility profile in the Set Profile field. Example: On the Table Spaces Utilities Menu, type the name of the profile you want in the Set Profile field.
2	Type the number of the Set Profile option in the Option field.
3	Press Enter. Result: The system displays the message indicating the utility profile is set.

Setting the active utility profile from a list of existing profiles

You can display a list of existing profiles from a utilities menu and select the active profile from the list. Follow these steps.

Step	Action
1	On the utilities menu for an object list panel, type the number of the profile option in the Option field.
2	Press Enter. Result: The system displays the Utility Profiles menu, which includes a list of the existing profiles.
3	Type S in the Select field for the profile you want.
4	Press Enter. Result: The system displays a menu listing options for creating profiles.
5	Use the END command. Result: The system displays the Utility Profiles menu.
6	Use the END command again. Result: The system redisplay the utilities menu for the object list with the name of the profile you selected in the Set Profile field.

Determining what to do next

Once you have set the utility profile you want to use, see “Accessing Object List Panels to Use with Utilities” on page 344.

Accessing the Utilities

Overview

This unit contains background and instructions for accessing a utility.

Background about filtering and control of objects

You can use filtering to restrict the objects listed on a panel and, therefore, to restrict the objects that commands act upon. If you need a reminder about how to use filtering, see the chapter “Filtering Object Lists” in *!DB/WORKBENCH Basic User's Guide*. If you want to review all the methods for filtering available for object list panels, see the chapter “Methods for Filters—Object List Panels” in *!DB/WORKBENCH Reference Manual*.

Background about using the Util (U) select to reduce the number of objects

You can reduce the number of objects by using the Util (U) select instead of the UTIL command. With the Util (U) select, the utility affects only the single table space, index, or storage group you select. Using the Util (U) select results in the display of the menu for the utilities for the object list.

Background about parameters for the UTIL command

You can use the name of the utility as a parameter for the UTIL command. In this way, you can access the panels required for a specific utility without going through the utilities menu. For example, on the Table Spaces panel, you access panels required for the COPY utility by using this command on the command line.

UTIL COPY

Accessing a utility from the command line

Follow these steps to access a utility directly from the command line without displaying a utilities menu.

Step	Action
1	On the object list panel, type UTIL followed by a space and the name of the utility. Example: If you want to use the REORG utility, type UTIL REORG on the command line.
2	Press Enter. Result: The system displays a panel with fields for entry for the utility whose name you entered.

Accessing a utility from a utilities menu

Follow these steps to access a utility from a utilities menu.

Step	Action
1	On the object list panel that supports the utility you want, perform the appropriate action: <ul style="list-style-type: none"> ● For a single object, type U in the Select field ● For all objects in the list, type UTIL on the command line
2	Press Enter. Result: The system displays the utilities menu for the object list panel.
3	Type the number associated with the utility in the Option field.
4	Press Enter. Result: The system displays a panel with fields for entry for the utility you selected.

Determining what to do next

For information about how to complete values required to use the utility, see “Completing Values for the Utility” on page 356.

Completing Values for the Utility

Overview

This unit contains background and instructions for methods to complete values required to use the utility.

Background about results after all the panels are completed

Many of the utilities display a Process Member menu at the end of the process. Some utilities, such as the DIAG ABENDS utility, take effect once you press Enter.

Variations in the panels displayed when you use a utility

The panels that !DB/WORKBENCH displays when you use a utility can vary because of these conditions:

- You select a utility that supports partitions and you select one or more partitioned objects. (!DB/WORKBENCH automatically displays the Partition Specifications panel when you use the utility.)
 - You use values from one utility that means that you also want to use additional utilities. !DB/WORKBENCH automatically displays all of the panels for the supporting utilities that are associated with the utility you have chosen to run. If you want to bypass any or all of these panels, use the GEN command as discussed later in this unit.
-

Example of variations based on values you use

The values you use also affects the panels that !DB/WORKBENCH automatically displays. The chart shows an example of the differences in panels that !DB/WORKBENCH displays depending on the values you use for COPY and RUNSTATS for the REORG utility.

Value Used on the REORG Panel	Panels that !DB/WORKBENCH Displays
COPY=Y RUNSTATS=Y	5 panels (2 panels specific to the REORG utility, 1 for the COPY dataset, 1 for the RECOVERY dataset, and the General Options panel)
COPY=N RUNSTATS=N	3 panels (2 panels specific to the REORG utility, and the General Options panel)

Options for completing values online

You have two choices for completing the values required for using the utility:

- Complete the fields on each panel. Press Enter each time you complete a panel to display the next panel in the sequence. The last panel in the sequence is the General Options panel.
- Accept the values for the utility that are defined in your profile. To do so, bypass all of the panels for the utility and go directly to the General Options panel by using the GEN command.

Using the GEN command

You can use the GEN command when using the utility definition process online. You cannot use GEN in macros or in batch processing. To use the GEN command, enter it on the command line of any utility definition panel and press Enter.

Example of using the GEN command

This chart provides an example of the use of the GEN command.

Step	Action
1	On the Table Spaces panel, type UTIL COPY on the command line.
2	Press Enter. Result: The system displays the COPY Utility Options panel (the first panel for which you can enter values for the COPY utility).
3	On the COPY Utility Options panel, type GEN on the command line.
4	Press Enter. Result: !DB/WORKBENCH uses values from your profile to complete the Copy Options and immediately displays the General Options panel (the last panel of those required for the Copy utility).

Reviewing your entries for accuracy

When you have completed your data entry on a utility profile panel, always press ENTER. Doing so causes !DB/WORKBENCH to check your entries for accuracy. If !DB/WORKBENCH discovers any errors, you must manually correct them. !DB/WORKBENCH does not automatically perform this function.

Avoid pressing END to exit from a data entry panel, since this will save your entries but will *not* check their accuracy.

Determining what to do next

For instructions for processing the utility statements, see “Processing the Statements Generated” on page 359.

Processing the Statements Generated

Overview

This unit contains background about the statements generated, the options available for processing, and the procedures to follow if the utility you use displays the Process Member menu.

Background about options from the Process Member menu

Most utilities display a Process Member menu after you complete values on the panels required for the utility. The menu provides options for statements generated by the utilities and by !DB/WORKBENCH. The chart lists the options and the response of the system when you select the option.

Options on the Menu	System Response
SUBMIT	Submits the statements generated
EDIT	Invokes the ISPF Editor and displays the contents of the member (When you complete editing, the system redisplay the Process Member menu.)
DELETE	Deletes the member from the !DB/WORKBENCH user PDS; returns to the Utilities Menu for the object list panel
PRINT	Prints the contents of the member to a JES printer
END	Leaves the member in the !DB/WORKBENCH user PDS and redisplay the Utilities Menu (The next time you use a utility, the system overwrites the contents in the member.)

Background about generated statements

When you use a DB2 utility or other options available from a menu of utilities available within !DB/WORKBENCH, !DB/WORKBENCH generates these items:

- JOB Card
- Utility EXEC and DD statements that are tailored according to your installation
- Utility control statements for the specific utility

Background about generation of unique DD and dataset names

If you use the standard options for the DB2 utilities, !DB/WORKBENCH generates unique names. However, some options for the DB2 utilities can generate duplicate names that cause errors. If a potential problem exists, !DB/WORKBENCH displays a warning message when you take any of the following actions:

- Type **1** in the JOB CONTROL field of the General Options panel
- Type a specific name for a DDname
- Use the object (space) name for a DDname when working with partitioned objects

Procedure for processing the statements

Follow these steps to process the statements.

Step	Action
1	On the Process Member menu, type the number of the option in the Option field.
2	Press Enter. Result: The system processes the statements according to the option you selected.

Accumulating Statements from Multiple Utilities

Overview

This unit provides background and instructions for accumulating the generated statements for several utilities.

Background about accumulating statements

You can use !DB/WORKBENCH to execute jobstreams from a combination of DB2 utilities. You accumulate statements in the process member for the utilities by using the MOD parameter with the UTIL command. Each time you use the parameter, the system appends any existing statements. When you have the combination that you want, you can then execute the accumulated statements from the Process Member menu.

Background about parameters for the UTIL command

You can add the MOD parameter to UTIL *utilityname* to append the current statements in the process member for utilities. For example, you can append the generated statements from the COPY utility to the existing statements in the process member by using this command on the command line.

```
UTIL COPY MOD
```

Implications for batch processing

If you want to use the MOD parameter in batch these additional resources exist:

- “EXEC Command” on page 554 and “NOEXEC Command and the Batch Utility” on page 558
- “UTIL Command and Functions Supported in Batch” on page 572

Characteristics of the example

In the example, you want to use utilities with the table space TEST96 and you have previously filtered the list. You want to accumulate statements from these utilities and execute them in the same jobstream:

- COPY
- REORG

Example of accumulating statements online

Follow these steps.

Step	Action
1	On the Table Spaces panel, type UTIL COPY on the command line.
2	Press Enter. Result: The system displays the first panel for the COPY utility.
3	Complete all panels required for the COPY utility.
4	Press Enter. Result: The system displays the Process Member menu. (The member contains statements from the COPY utility.)
5	Use the END command. Result: The system displays the Table Spaces panel.
6	On the Table Spaces panel, type UTIL REORG MOD on the command line.
7	Press Enter. Result: The system displays the Process Member menu. (The member contains statements from the COPY utility and from the REORG utility.)
8	On the Process Member menu, type 1 in the Option field.
9	Press Enter. Result: !DB/WORKBENCH submits the statements for processing.

Chapter 16.

DB2 Utilities and Options Available with !DB/WORKBENCH

Introduction

This chapter provides an introduction to DB2 utilities and options in !DB/WORKBENCH. The chapter lists the utilities in alphabetical order and provides prerequisites and any required background for each of the utilities.

Chapter contents

Overview of the Chapter	364
Online Help and Specific Values for Utilities	366
CHECK DATA Utility	367
CHECK INDEX Utility	369
COPY Utility	371
DIAG ABENDS Utility	375
DIAG SYSUTIL Utility	377
DISASTER REC Utility	379
ICCAT Utility	382
MERGE COPY Utility	384
MODIFY Utility	387
QUIESCE Utility	389
REBIND Command	391
REBIND PKGS Command	394
REBIND PLANS Command	397
RECOVER INDEX Utility	400
RECOVER Utility for Table Spaces	402
REORG INDEX Utility	407
REORG Utility for Table Spaces	409
REPORT Utility	411
RUNSTATS Utility	413
STOSPACE Utility	417

Overview of the Chapter

Prerequisites for using this chapter

Before using any of the utilities in this chapter, review the information in “Introduction to DB2 Utilities and Options” on page 339 and “Panels and Fields Used with Multiple Utilities” on page 419.

Organization of information in the chapter

This chapter provides information specific to each DB2 utility and option available with DB2. The utilities are listed in alphabetical order. The chapter provides this information about each of the DB2 utilities and options available with !DB/WORKBENCH:

- Introduction
 - Availability on object list panels (including with RI data if available)
 - Authority required to use the utility
 - Background about the utility (if any)
 - Specific values you must provide to use the utility
 - Sources of information about values that are used by multiple utilities
-

Background about these fields and panels and the utility profiles

You can begin to use DB2 utilities and other options without defining values for many of the fields covered in this chapter. !DB/WORKBENCH provides a default utility profile that it uses if you do not create your own.

Relationship of this chapter to other chapters in this section

This chapter is intended to be used in conjunction with the next chapter, “Panels and Fields Used with Multiple Utilities” on page 419. That chapter covers information common to most of the utilities and options available with !DB/WORKBENCH. This chapter covers information specific to most of the utilities.

For example, you can review prerequisites for using the CHECKDATA utility in this chapter. Depending on your experience with other !DB/WORKBENCH utilities or your satisfaction with your current utility profile, you can decide whether you also want to review prerequisites covered in “Options for DB2 Commands with Utilities” on page 426, “General Options for Job Control” on page 428, “Values for Identifying and Restarting a Utility” on page 431, and “Values for the Work Data Sets” on page 432.

Utilities available for both table spaces and indexes and this chapter

Several utilities are available on both the table spaces and indexes panel. When there are no differences in operation or minimum differences, the information is presented in one unit. When the differences are substantial, the utility is discussed separately for the table space and index. This occurs in these cases:

- RECOVER utility
- REORG utility

Additional resources for using utilities efficiently

These additional resources are available for creating and setting utility profiles.

Task You Want to Perform	Source to Use
Customize one or more utility profiles	“Creating or Modifying Profiles for DB2 Utilities and Options” on page 449
Use a utility with partitioned objects	“Values for Partitions” on page 442

Methods for identifying additional fields for utilities

Several methods are included in this section to help you determine quickly what information common to most utilities (and covered in “Panels and Fields Used with Multiple Utilities” on page 419) is appropriate for the specific utility you want to use.

- The sources of information included with each utility in this chapter
- The chart included in the unit “Overview of the Chapter” on page 420 in the chapter “Panels and Fields Used with Multiple Utilities” on page 419.
- The introduction to each unit in “Panels and Fields Used with Multiple Utilities” on page 419. For example, the information about “General Options for Job Control” on page 428 begins with a list of each utility that uses the options.

Online Help and Specific Values for Utilities

Overview

This chapter provides prerequisites and any required background for each of the utilities. The online Help provides detailed information about values for each field.

Help and determining the correct entry for fields for utilities

If you want to know more information about what value is appropriate for the way you want to use the utility, type **HELP** on the command line and press Enter, or press the appropriate function key. The online Help panel is typical of the resources available.

```

Cmd ==>
----- HELP FOR CHECK DATA UTILITY -----
Press END to return, ENTER to continue
or TAB to highlighted text and press ENTER to display more information.

Description      Use this utility to check the integrity of referential
                  tables. The Check Data utility checks that every foreign
                  key value matches a primary key value or is null.

UTILITY-ID      Type the utility ID you want assigned to this utility. The
                  default ID is your 'userid.jobname', but you can specify any
                  16 characters you want. The system sets the Utility-ID
                  variable (UID) in the JCL EXEC statement to the ID you type.

START           If you are not restarting a faulty utility, leave blank.

                  If you want to restart the utility at the last commit point,
                  type C.

                  If you want to restart the utility at the beginning of the
                  last phase executed, type P.

-----
GLOSSARY  INDEX  COMMANDS  FILTERS  WINDOWS  SELECTS  MORE->

```

CHECK DATA Utility

Overview

This unit introduces the CHECK DATA utility. The unit covers availability and background about the CHECK DATA utility. It also contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the CHECK DATA utility

The CHECK DATA utility checks the integrity of referential tables. The utility checks that every foreign key value matches a primary key value or is null. The utility also allows you to update or use an existing exception table, and then it deletes rows that violate referential constraints and copies them to the exception table.

Availability of the CHECK DATA utility

The CHECK DATA utility is available from the Table Spaces panel; however, to use it with referential data, you need to first access (either online or in batch) panels with referential data. You have two methods available from the Tables panel:

- Use the TSSET command; then use the UTIL command.
 - Use the RI Info (F) select with a table with referential data; then use the SET command; then use the UTIL command.
-

Authority required to use the CHECK DATA utility

If you want to use the option for the CHECK DATA utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
 - DBADM, DBCTRL, or DBMAINT authority
 - STATS privilege for the database
-

Background about the CHECK DATA utility

The CHECK DATA utility affects table space sets (table spaces of referentially related tables). If you use the CHECK DATA utility on a panel with a table space set, then the system generates CHECK DATA utility statements in sets. If you specify a single job step (JOBCONTROL = 1 on the General Options panel), then the system checks the integrity of referential tables for each set in a different step.

Values to provide that are used by multiple utilities

To use the CHECK DATA utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for DB2 commands (Values for CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart	“Values for Identifying and Restarting a Utility” on page 431
Values for the utility dataset	“Values for the Work Data Sets” on page 432
Values for the sort work datasets	“Values for Sort Work Data Sets” on page 435

Specific values to provide when you use the utility

When you use the CHECK DATA utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
SCOPE	P — Do only marked tables A — Do all tables
DELETE	Y/N — Remove invalid rows and descendents?
COPY	Y/N — Copy invalid rows to Exception Table?
MAX ERRORS	Max # Errors to tolerate. Stop after this #.

CHECK INDEX Utility

Overview

This unit introduces the CHECK INDEX utility. It covers availability of the CHECK INDEX utility and also contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the CHECK INDEX utility

The CHECK INDEX utility tests whether indexes are consistent with the data they index.

- On the Table Spaces panel, the utility checks all indexes within the table space.
 - On the Indexes panel, the utility checks index-to-table pointers.
-

Availability of the CHECK INDEX utility

The CHECK INDEX utility is available from these object list panels:

- Table Spaces panel
 - Indexes panel
-

Authority required to use the CHECK INDEX utility

If you want to use the option for the CHECK INDEX utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM, DBCTRL, or DBMAINT authority for the database
- STATS privilege for the database

Values to provide that are used by multiple utilities

To use the CHECK INDEX utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for DB2 commands (Values for CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart	“Values for Identifying and Restarting a Utility” on page 431
Values for the utility datasets	“Values for the Work Data Sets” on page 432
Values for the sort work datasets	“Values for Sort Work Data Sets” on page 435

COPY Utility

Overview

This unit introduces the COPY utility. It covers availability of the COPY utility and contains information about values you provide when you use the utility. It includes information about image copies within !DB/WORKBENCH. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the COPY utility

The COPY utility writes an image copy of a table space by using the DB2 COPY command.

Availability of the COPY utility

The COPY utility is available from the Table Spaces panel.

Authority required to use the COPY utility

If you want to use the option for the COPY utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM, DBCTRL, or DBMAINT authority for the database
- IMAGCOPY privilege for the database

Types of image copies available with !DB/WORKBENCH

You can specify various options for the IMAGE COPY field as shown in this chart:

Field Name	Description on the Panel
IMAGE COPY	<p>Y Perform a full image copy</p> <p>N Perform an incremental image copy</p> <p>M Perform an incremental image copy with:</p> <ul style="list-style-type: none"> ● Only pages changed since the last image copy ● Accurate page sizes. !DB/WORKBENCH uses the DB2 space map pages of each dataset to determine the exact number of pages updated. <p>C Full image copy using the DFSMS concurrent copy facility</p>

DB2 space map pages and consequences when using IMAGE COPY=M

When !DB/WORKBENCH makes the incremental copy, it uses the DB2 space map pages of each dataset to determine the number of pages. Because this method reads the physical VSAM space (and does not interact with DB2), using IMAGE COPY=M results in these consequences on the sizing of objects:

- Includes the exact number of pages based on the DB2 space mMap pages
- Does not include pages in the DB2 buffer pools that DB2 has not yet externalized

For example, if you have made changes but DB2 has not yet externalized these changes and if you use IMAGE COPY=M, the system does not recognize these pages as being updated and does not include the pages in the sizing of the objects for the incremental copy. To avoid having table spaces that are changed but are not included in the image copies, you can use the QUIESCE utility with the option WRITE YES (DB2 V3.1 or higher) before you use the COPY utility. (The option WRITE YES with the QUIESCE utility externalizes all pages that have changed.)

Implications for using IMAGECOPY=M

These implications exist in using IMAGECOPY=M:

Condition with IMAGE COPY=M	Implication
You use IMAGE COPY=M with 10 objects. (6 have been modified; 4 are unchanged)	!DB/WORKBENCH acts only on the 6 objects that have been modified and excludes the 4 that were unchanged. (Incremental copies consist only of pages that have been changed.)
You use IMAGE COPY=M and also use MIGRATED=I (on the General Options panel) because you want to include a migrated object in the incremental image copy.	If you include migrated objects, !DB/WORKBENCH ignores the IMAGE COPY=M value and instead uses the allocated size that is set by DFHSM.

Requirements for using IMAGE COPY=M

Because the system uses the DB2 space map page with IMAGE COPY=M, you must have authorization to OPEN and READ the VSAM datasets.

Values to provide that are used by multiple utilities

To use the COPY utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for DB2 commands (Values for CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart	“Values for Identifying and Restarting a Utility” on page 431
Values for copy datasets for the local site	“Values for Copy Data Sets (LOCAL)” on page 438
Values for copy datasets for the recovery site	“Values for Copy Data Sets (RECOVERY)” on page 440

Specific values to provide when you use the utility

When you use the COPY utility, you provide the values for the fields in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
SHRLEVEL	R for REFERENCE or C for CHANGE
IMAGE COPY	Y for full or N for Incremental Image Copy M to read space map pages for accurate Incremental Image Copy sizing C for a full image copy using the DFSMS concurrent copy facility

DIAG ABENDS Utility

Overview

This unit introduces the DIAG ABENDS utility. It covers availability of the DIAG ABENDS utility and contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the DIAG ABENDS utility

The DIAG ABENDS utility generates DIAGNOSE control statements and other utility control statements for utilities. These statements generate information useful in diagnosing problems.

Availability of the DIAG ABENDS utility

The DIAG ABENDS utility is available from these object list panels:

- Table Spaces panel
 - Indexes panel
-

Authority required to use the DIAG ABENDS utility

If you want to use the option for the DIAG ABENDS utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
 - DBADM or DBCTRL authority for the database
 - REPAIR privilege for the database
-

Background about the DIAG ABENDS utility

Unlike most of the other DB2 utilities, the DIAG ABENDS utility does not display the Process Member menu for JCL statements. It redisplay the UTIL selection list.

Specific values to provide when you use the utility

When you use the DIAG ABENDS utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
ALLDUMPS	Y/N — Force a dump for all utility ABENDs? to specify one or more ABEND conditions
ABENDS	Y/N — Set specific ABEND conditions If Y, additional screens will be presented to identify one or more ABEND conditions.

Additional specific values required

The value you specify for the ABENDS field determines any additional values required. (The values in the chart can be specified in the profile.)

Condition	Action to Perform
If on the first panel you specified Y for ABENDS field	Type the value in the field for one of these conditions: MESSAGE Message number, that is U521 (drop the “DSN” prefix). TRACEID Trace Identifier (four hex digits)
If on the second panel a value in the MESSAGE or TRACEID field	INSTANCE How many times MESSAGE or TRACEID occurs before an ABEND dump is to be taken. Specify a number. NODUMP Y/N Suppress the Dump for this ABEND. This option overrides ALLDUMP for this MESSAGE or TRACEID condition.

DIAG SYSUTIL Utility

Overview

This unit introduces the DIAG SYSUTIL utility. It covers availability of the DIAG SYSUTIL utility and contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the DIAG SYSUTIL utility

The DIAG SYSUTIL utility lists SYSUTIL rows and the DB2 Directory contents for a table space or index space.

Availability of the DIAG SYSUTIL utility

The DIAG SYSUTIL utility is available from the Table Spaces panel.

Authority required to use the DIAG SYSUTIL utility

If you want to use the option for the DIAG SYSUTIL utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM or DBCTRL authority for the database
- REPAIR privilege for the database

Values to provide that are used by multiple utilities

To use the DIAG SYSUTIL utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for job control	“General Options for Job Control” on page 428

Specific values to provide when you use the utility

When you use the DIAG SYSUTIL utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
SYSUTIL	Y/N — Display all SYSUTIL rows
DIRECTORY	Displays DB2 Directory contents for Table Spaces One of: N Do not display A Display ALL object entries T Display only Table object entries I Display only Index object entries

DISASTER REC Utility

Overview

This unit introduces the DISASTER REC utility. This unit covers availability and background about the DISASTER REC utility. It also contains information about the values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the DISASTER REC utility

The DISASTER REC utility is a !DB/WORKBENCH utility that you can use to recover a tablespace. When you use the DISASTER REC utility, !DB/WORKBENCH:

- creates the table space using the DDL in the extract you specified for your session
- creates the dependent objects you specify using the DDL in the extract you specified for your session
- recovers the data for the table space using the last image copy that was saved

Availability of the DISASTER REC utility

The DISASTER REC utility is available from the Table Spaces panel. You can also use the utility with all the table spaces for the database you select using the S (Table spaces) select on the Databases object list panel.

Authority required to use the DISASTER REC utility

If you want to use the option for the DISASTER REC utility, you must have these authorizations:

- authority to create the table space
- authority to create the dependent objects you specify for the table space
- read and write authority for the VSAM datasets for the table space
- read authority for the dataset that contains the image copy

Additional information about the DISASTER REC utility

When you use the DISASTER REC utility to recover a table space, you can also specify that !DB/WORKBENCH include other utilities and commands in the process for recovering the table space. For example, you can specify that !DB/WORKBENCH include the CHECK DATA utility in the process for recovering the table space.

Values to provide that are used by multiple utilities

To use the DISASTER REC utility, you provide values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for DB2 commands (Values for CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart	“Values for Identifying and Restarting a Utility” on page 431
Values for copy datasets for the local site	“Values for Copy Data Sets (LOCAL)” on page 438
Values for copy datasets for the recovery site	“Values for Copy Data Sets (RECOVERY)” on page 440

Specific values to provide when you use the utility

When you use the DISASTER REC utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
CHECK PAGES	Y Check input data pages N Skip check processing
RESET LOGRBA	Y Reset N Skip Reset
RECOVER INDEXES	Y Include Indexes N Skip Indexes
INCLUDE RELATED	Y Include all related objects N Skip related obj
NEW IMAGE COPY	Y/N/F Create/Force a new image copy after recovery
CHEKC DATA	Y/N Generate CHECK DATA on Referential Sets
RUNSTATS	Y/N Generate RUNSTATS before any REBINDs
REBIND	Y/N REBIND plans and packages

ICCAT Utility

Overview

This unit introduces the ICCAT utility. It covers availability of the ICCAT utility and contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the ICCAT utility

The ICCAT utility catalogs image copy datasets.

Availability of the ICCAT utility

The ICCAT utility is available from the Table Spaces panel.

Authority required to use the ICCAT utility

If you want to use the option for the ICCAT utility, you must have one of these authorities:

- SYSADM or SYSCTRL authority
 - DBADM, DBCTRL, or DBMAINT authority for the database
 - IMAGCOPY privilege for the database
-

Values to provide that are used by multiple utilities

To use the ICCAT utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for job control	“General Options for Job Control” on page 428

Specific values to provide when you use the utility

There are no values specific to the ICCAT utility. When you use the utility, the utility displays the General Options panel.

If you want to specify datasets to be migrated

If you want to specify which datasets to migrate, do not provide a value in the Migrated field. When you press Enter, the system displays a panel listing the name of each dataset in turn. You can either skip each dataset or include it.

MERGE COPY Utility

Overview

This unit introduces the MERGE COPY utility. It covers availability of the MERGE COPY utility and contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the MERGE COPY utility

The MERGE COPY utility combines one or more incremental image copies of a table space to make one incremental image copy or merges one or more incremental image copies with a full image copy to make a new full image copy.

Availability of the MERGE COPY utility

The MERGE COPY utility is available from the Table Spaces panel.

Authority required to use the MERGE COPY utility

If you want to use the option for the MERGE COPY utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM, DBCTRL, or DBMAINT authority
- IMAGCOPY privilege for the database

Values to provide that are used by multiple utilities

To use the MERGE COPY utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for DB2 commands (Values for CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart	“Values for Identifying and Restarting a Utility” on page 431
Values for the utility datasets	“Values for the Work Data Sets” on page 432
Values for copy datasets for the local site	“Values for Copy Data Sets (LOCAL)” on page 438
Values for copy datasets for the recovery site	“Values for Copy Data Sets (RECOVERY)” on page 440

Specific values to provide when you use the utility

When you use the MERGE COPY utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
NEWCOPY	Y — Merge with full copy to form new full copy N — Merge only incremental copies
WORKDDNM	# = WORK <i>nnnn</i> default = SYSUT1
WORKUNIT	Device type for copy output datasets
WORKSIZE	Override copy output dataset size (tracks)

MODIFY Utility

Overview

This unit introduces the Modify Recovery utility. It covers availability of the Modify Recovery utility and contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the MODIFY utility

The Modify Recovery utility deletes unwanted image copy records from the DB2 SYSCOPY table and the related log records from the DB2 log.

Availability of the MODIFY utility

The MODIFY utility is available from the Table Spaces panel.

Authority required to use the MODIFY utility

If you want to use the option for the MODIFY utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM, DBCTRL, or DBMAINT authority for the database
- IMAGCOPY privilege for the database to run MODIFY RECOVERY
- STATS privilege for the database to run MODIFY STATISTICS

Values to provide that are used by multiple utilities

To use the MODIFY utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart	“Values for Identifying and Restarting a Utility” on page 431

Specific values to provide when you use the utility

When you use the MODIFY RECOVERY utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
AGE	Delete all records older than AGE days. (0=*)
DATE	Delete all records older than date <i>yyyymmdd</i> . The date cannot be a date in the future.

QUIESCE Utility

Overview

This unit introduces the QUIESCE utility. The unit covers availability and background about the QUIESCE utility. It also contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the QUIESCE utility

The QUIESCE utility stops further changes from being made to the DB2 database and records the log RBA of the quiesce point for recovery of data.

Availability of the QUIESCE utility

The QUIESCE utility is available from the Table Spaces panel.

Authority required to use the QUIESCE utility

If you want to use the option for the QUIESCE utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
 - DBADM, DBCTRL, or DBMAINT authority for the database
 - IMAGCOPY privilege for the database
-

Background about QUIESCE and the JOBCONTROL value

Review the chart for two methods to use the QUIESCE utility and how your choice affects the value you provide for JOBCONTROL on the General Options panel.

Way You Want to Use the Quiesce Utility	Action to Perform
With each object having a separate recovery point	On the General Options panel, use JOBCONTROL=2
With all objects listed having a single recovery point (typically used with RI data)	On the General Options panel, use JOBCONTROL=1

Prerequisites for a single recovery point

If you are using the QUIESCE utility to cause all objects listed to have a single recovery point, one restriction applies. The maximum number of objects in one QUIESCE is 255.

Specific values to provide when you use the utility

To use the QUIESCE utility, you provide the value for the field shown in the chart.

Field Name	Description on Panel
WRITE	Y/N — Externalize pages?

Values to provide that are used by multiple utilities

To use the QUIESCE utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Type of Values	Source to Use
Options for DB2 commands (Values for CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart	“Values for Identifying and Restarting a Utility” on page 431

REBIND Command

Overview

This unit introduces the REBIND command, which is available from the Utilities Menu. The unit covers availability and background about the REBIND command. It also contains information about values you provide when you use the command. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Relationship of the REBIND command to REBIND PLANS and REBIND PKGS

Use the REBIND command when you wish to rebind both plans and packages. You can selectively rebind either plans or packages by using the REBIND PLANS command or REBIND PKGS command respectively. See “REBIND PKGS Command” on page 394 and “REBIND PLANS Command” on page 397 later in this chapter for additional information.

Introduction to the REBIND command

The REBIND command rebinds the plans and packages you selected. !DB/WORKBENCH generates the JCL required to run the DB2 commands in batch.

Availability of the REBIND command

The REBIND command is available from these object list panels:

- Table Spaces panel
- Indexes panel

Authority required to use the REBIND command

If you want to use the option for the REBIND command, you must have one of the authorizations for the plan or one of these authorizations for the package:

- SYSADM or SYSCTRL authority
- BIND privilege
- BINDAGENT privilege
- PACKADM authority (packages only)
- Ownership of the plan (plans only)

Values to provide that are used by multiple utilities

To use the REBIND command, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for job control	“General Options for Job Control” on page 428

Specific values to provide when you use the command

When you use the REBIND command, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panels.)

Field Name	Description on the Panel
Flag	I Informational W Warning E Error C Completion
Validate	R Run B Bind
Isolation	R RR C CS U UR
Acquire	U Use A Allocate
Release	C Commit D Deallocate
Deferprep	Y Yes N No
Explain	Y Yes N No

Specific values to provide when you use the command (continued)

The chart contains the remaining fields for the REBIND command.

Field Name	Description on the Panel
SqlRules	D DB2 S Standard
Dynamicrules	R Run B Bind blank
Degree	1, ANY or Blank 1 prohibits I/O parallelism
Disconnect	E Explicit A Automatic C Conditional
Cachesize	0, 256-4096 (0 inhibits caching, default 1024)
Qualifier	Qualifier of all unqualified object names
Currentdata	Y Yes N No
Currentserver	No description (Default = location name of the current server)
Menu option 1	Use current owner
Menu option 2	Use Plan creator as Owner
Menu option 3	Use User ID
Menu option 4	Use Override Owner
Batch	Enable E or Disable D
DLIbatch	Enable E or Disable D
DB2call	Enable E or Disable D
CICS	Enable E or Disable D
IMS	Enable E or Disable D
IMSBMP	Enable E or Disable D
IMSMPP	Enable E or Disable D
Remote	Enable E or Disable D

REBIND PKGS Command

Overview

This unit introduces the REBIND PKGS command, which is available from the Utilities Menu. The unit covers availability and background about the REBIND PKGS command. It also contains information about values you provide when you use the command. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Relationship of the REBIND PKGS command to the REBIND command

Use the REBIND command when you wish to rebind both plans and packages. You can selectively rebind either plans or packages by using the REBIND PLANS command or REBIND PKGS command respectively. See “REBIND Command” on page 391 and “REBIND PLANS Command” on page 397 in this chapter for additional information.

Introduction to the REBIND PKGS command

The REBIND PKGS command rebinds the packages you selected. !DB/WORKBENCH generates the JCL required to run the DB2 commands in batch.

Availability of the REBIND PKGS command

The REBIND PKGS command is available from these object list panels:

- Table Spaces panel
 - Indexes panel
-

Authority required to use the REBIND PKGS command

If you want to use the option for the REBIND PKGS command, you must have one of the authorizations for the plan or one of these authorizations for the package:

- SYSADM or SYSCTRL authority
 - BIND privilege
 - BINDAGENT privilege
 - PACKADM authority (packages only)
 - Ownership of the plan (plans only)
-

Values to provide that are used by multiple utilities

To use the REBIND PKGS command, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for job control	“General Options for Job Control” on page 428

Specific values to provide when you use the command

When you use the REBIND PKGS command, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panels.)

Field Name	Description on the Panel
Flag	I Informational W Warning E Error C Completion
Validate	R Run B Bind
Isolation	R RR C CS N NC S RS U UR
Release	C Commit D Deallocate

Specific values to provide when you use the command (continued)

The chart contains the remaining fields for the REBIND PKGS command.

Field Name	Description on the Panel
Explain	Y Yes N No
Dynamicrules	R Run B Bind Blank
Degree	1, ANY or Blank 1 prohibits I/O parallelism
Currentdata	Y Yes N No
Deferprep	Y Yes N No
Reoptimize	Y Yes N No
Menu option 1	Use current owner
Menu option 2	Use Plan creator as Owner
Menu option 3	Use User ID
Menu option 4	Use Override Owner
Batch	Enable E or Disable D
DLIbatch	Enable E or Disable D
DB2call	Enable E or Disable D
CICS	Enable E or Disable D
IMS	Enable E or Disable D
IMSBMP	Enable E or Disable D
IMSMPP	Enable E or Disable D
Remote	Enable E or Disable D

REBIND PLANS Command

Overview

This unit introduces the REBIND PLANS command, which is available from the Utilities Menu. The unit covers availability and background about the REBIND PKGS command. It also contains information about values you provide when you use the command. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Relationship of the REBIND PLANS command to the REBIND command

Use the REBIND command when you wish to rebind both plans and packages. You can selectively rebind either plans or packages by using the REBIND PLANS command or REBIND PKGS command respectively. See “REBIND Command” on page 391 and “REBIND PKGS Command” on page 394 in this chapter for additional information.

Introduction to the REBIND PLANS command

The REBIND PLANS command rebinds the packages you selected. !DB/WORKBENCH generates the JCL required to run the DB2 commands in batch.

Availability of the REBIND PLANS command

The REBIND PLANS command is available from these object list panels:

- Table Spaces panel
- Indexes panel

Authority required to use the REBIND PLANS command

If you want to use the option for the REBIND PLANS command, you must have one of the authorizations for the plan or one of these authorizations for the package:

- SYSADM or SYSCTRL authority
- BIND privilege
- BINDAGENT privilege
- PACKADM authority (packages only)
- Ownership of the plan (plans only)

Values to provide that are used by multiple utilities

To use the REBIND PLANS command, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for job control	“General Options for Job Control” on page 428

Specific values to provide when you use the command

When you use the REBIND PLANS command, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panels.)

Field Name	Description on the Panel
Flag	I Informational W Warning E Error C Completion
Validate	R Run B Bind
Isolation	R RR C CS U UR
Acquire	U Use A Allocate
Release	C Commit D Deallocate
Deferprep	Y Yes N No
Reoptimize	Y Yes N No
Explain	Y Yes N No

Specific values to provide when you use the command (continued)

The chart contains the remaining fields for the REBIND PLANS command.

Field Name	Description on the Panel
SqlRules	D DB2 S Standard
Dynamicrules	R Run B Bind Blank
Degree	1, ANY or Blank 1 prohibits I/O parallelism
Disconnect	E Explicit A Automatic C Conditional
Cachesize	0, 256-4096 (0 inhibits caching, default 1024)
Qualifier	Qualifier of all unqualified object names
Currentdata	Y Yes N No
Currentserver	No description (Default = location name of the current server)
Menu option 1	Use current owner
Menu option 2	Use Plan creator as Owner
Menu option 3	Use User ID
Menu option 4	Use Override Owner
Batch	Enable E or Disable D
DLIbatch	Enable E or Disable D
DB2call	Enable E or Disable D
CICS	Enable E or Disable D
IMS	Enable E or Disable D
IMSBMP	Enable E or Disable D
IMSMPP	Enable E or Disable D
Remote	Enable E or Disable D

RECOVER INDEX Utility

Overview

This unit introduces the RECOVER INDEX utility for indexes. It covers availability of the RECOVER INDEX utility and contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the RECOVER INDEX utility

The RECOVER INDEX utility recreates indexes from the table that they reference.

Availability of the RECOVER INDEX utility

The RECOVER INDEX utility is available from the Indexes panel. (The RECOVER utility is also available from the Table Spaces panel; the utility used with table spaces is described in a separate unit.)

Authority required to use the RECOVER INDEX utility

If you want to use the option for the RECOVER INDEX utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
 - DBADM or DBCTRL authority for the database
 - RECOVERDB privilege for the database
-

Background for remote recovery

The facility for generating JCL for the RECOVER utility checks for the existence of physical objects. However, if an object is missing, the facility still generates the JCL. If an object is missing when the recovery is for an index, !DB/WORKBENCH causes the JCL to be generated only if you specify the work and sort datasets. For information on sizing for these datasets, see “How !DB/WORKBENCH Determines Data Set Sizes” on page 446.

Values to provide that are used by multiple utilities

To use the RECOVER INDEX utility, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for DB2 commands (Values for the CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart (if desired)	“Values for Identifying and Restarting a Utility” on page 431
Values for the utility dataset	“Values for the Work Data Sets” on page 432
Values for the sort work datasets	“Values for Sort Work Data Sets” on page 435

Specific values to provide when you use the utility

There are no values specific to the RECOVER INDEX utility. When you use the utility, the utility displays fields that are used by multiple utilities.

RECOVER Utility for Table Spaces

Overview

This unit introduces the RECOVER utility for table spaces. The unit covers availability and background about the RECOVER utility. It also contains information about values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the RECOVER utility for table spaces

The RECOVER utility recovers DB2 data from image copies of a table space and from later table space change log records.

Availability of the RECOVER utility for table spaces

The RECOVER utility is available from the Table Spaces panel. (It is also available from the Indexes panel; the utility used with indexes is described in a separate unit.)

Authority required to use the RECOVER utility for table spaces

If you want to use the option for the RECOVER utility for table spaces, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
 - DBADM or DBCTRL authority for the database
 - RECOVERDB privilege for the database
-

Background about the RECOVER utility for table spaces

When planning a recovery of multiple table spaces, remember the number of table spaces that can be recovered during the execution of the DB2 RECOVERY utility is limited. This limit typically depends on these issues:

- The number of image copies (full and partial) that are needed to perform the recovery
- Whether any table spaces are part of a referential set since all table spaces in a referential set must be recovered simultaneously.

For more information on this topic, see the *IBM Database 2 Command and Utility Reference*.

Background about the utility and recovery of table space sets

If a table space set (table spaces of referentially related tables) needs to be recovered, you must either recover all the table spaces to the common recovery point or risk inconsistencies among the tables involved in the referential relationship. Follow these steps.

Step	Action
1	On the Table Spaces panel, type T in the field for the damaged table space.
2	Press Enter. Result: The system displays the Tables panel with a list of all the tables in the table space.
3	On the Tables panel, type TSSET on the command line.
4	Press Enter. Result: The system tests all the table spaces in the table space set.
5	On the panel showing table space sets, type UTIL RECOVER on the command line. Result: The system displays the first option panel for the RECOVERY utility
6	If necessary, complete the values on the first panel.
7	Press Enter. Result: The system displays the second panel of options.
8	Type ? in the TORBA field.
9	If necessary, complete the other values on the second panel. Result: The system dynamically accesses the SYSIBM.SYSCOPY table in the DB2 catalog and lists all QUIESCE and IMAGE COPY entries for one of the table spaces in the set.
10	Type an S in the field for the object you want.
11	Press Enter. Result: The system generates the RECOVER jobstream. (Because you used the utility from the panel showing table space sets, the system automatically generates a table space set RECOVER.)

Background for remote recovery

The facility for generating JCL for the RECOVER utility checks for the existence of physical objects. However, if an object is missing, the facility still generates the JCL. If an object is missing when the recovery is for an index, !DB/QUICKCHANGE causes the JCL to be generated only if you specify the work and sort datasets.

Values to provide that are used by multiple utilities

To use RECOVER, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for DB2 commands (Values for the CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart (if desired)	“Values for Identifying and Restarting a Utility” on page 431

Specific values to provide when you use the utility for table spaces

When you use the RECOVER utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
RECOVER SET?	Y/N to Recover all listed table spaces as a set. Y generates a single RECOVER utility statement for multiple spaces; N generates a RECOVER utility statement for each tablespace. Both use the TORBA, TOLOGPOINT, or LOGONLY parameters.
SITE	L for local site; R for remote site; blank for current site.
INCLUDE IX	Y to recover all indexes for the table space, or N.

Specific values to provide when recovering a single page

This chart continues the information about specific values to provide when you use the utility for table spaces.

Field Name	Description on the Panel
PAGE	Number of page to be recovered (decimal digits, or hex digits between quotes preceded by "X").
CONTINUE	Y to continue if a previous Recovery failed. (Optional if recovering a single page.)

Specific values to provide when you use the utility and limit recovery

These values are options available to limit recovery. (The values are displayed on the second panel available with the RECOVERY utility.) You can specify one of these values or none.

Field Name	Description on the Panel
ERR-RANGE?	Y to recover pages with reported I/O errors.
TORBA or TOLOGPOINT	Hex string (Log Relative Byte) to stop recovery. Type ? for SYSCOPY lookup. Use TORBA when running with DB2 Version 3 or earlier; TOLOGPOINT when running with DB2 Version 4.
LOG ONLY	Y to use only the log records for recovery or N
TOCOPY	Copy DSName This names a specific Image Copy dataset. Enter with quotes if not prefixed with your USERID.

Additional values required

The value you specify for the TOCOPY field determines any additional required values. (The three values shown in the chart can be specified in the profile.)

Condition	Action to Perform
If the TOCOPY name is not unique in SYSCOPY	Type a value either for CATALOGED or for VOLSER and SEQ fields: CATALOGED? Y if cataloged, N if not. VOLSER Volume Serial number SEQ File sequence number on the volume

REORG INDEX Utility

Overview

This unit introduces the REORG INDEX utility. It covers availability of the REORG INDEX utility and contains information about the values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the REORG INDEX utility

When used with indexes, the REORG INDEX utility reorganizes an index space to improve access performance and reclaim fragmented space.

Availability of the REORG INDEX utility

The REORG utility is available from the Indexes panel. (The REORG utility is also available from the Table Spaces panel; the utility with table spaces is described in a separate unit.)

Authority required to use the REORG INDEX utility

If you want to use the option for the REORG utility for table spaces, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
 - DBADM or DBCTRL authority for the database
 - REORG privilege for the database
-

Values to provide that are used by multiple utilities

To use REORG INDEX, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for DB2 commands (Values for the CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428

Values to provide that are used by multiple utilities (continued)

This chart continues with sources of information about values that are used by multiple utilities.

Types of Values	Source to Use
Values to identify the utility and restart (if desired)	“Values for Identifying and Restarting a Utility” on page 431
Values for the utility dataset	“Values for the Work Data Sets” on page 432
Values for the sort work datasets	“Values for Sort Work Data Sets” on page 435

Specific values to provide when you use the utility for indexes

When you use the REORG INDEX utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
AFTER UNLD	C — Continue P — Pause O — Unload only
RUNSTATS	Y/N to perform a RUNSTATS after the REORG
REBIND	Y/N to bind all related PLANS/PACKAGES

REORG Utility for Table Spaces

Overview

This unit introduces the REORG utility for table spaces. It covers availability of the REORG utility for table spaces and contains information about the values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the REORG utility

When used with table spaces, the REORG utility rebuilds a table space and all indexes on any table in the table space to redistribute free space and reestablish clustering.

Availability of the REORG utility for table spaces

The REORG utility is available from the Table Spaces panel. (It is also available from the Indexes panel; the utility used with indexes is discussed in a separate unit.)

Authority required for the REORG utility for table spaces

If you want to use the option for the REORG utility for table spaces, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM or DBCTRL authority for the database
- REORG privilege for the database

Important

If you are using compressed table spaces, the size of the work data sets should be based on recent RUNSTATS. If you do not have RUNSTATS, the size of the work dataset will be based on used space values, and if your RUNSTATS are not recent, the size of the work datasets may be based on an incorrect compression factor.

Values to provide that are used by multiple utilities

To use REORG, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for DB2 commands (Values for the CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart (if desired)	“Values for Identifying and Restarting a Utility” on page 431
Values for the utility dataset	“Values for the Work Data Sets” on page 432

Specific values to provide when you use the utility

When you use the REORG utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
RELOAD LOG	Y/N Log during Reload? N is faster.
SORT DATA	Y/N Perform an external sort.
AFTER UNLD	C (CONTINUE), P (PAUSE) or O (UNLOAD ONLY).
KEEPDICTIONARY	Y/N To reuse compression dictionary.
Image COPY	Y/N Generate Full Image Copy after Reorg.
RUNSTATS	Y— Perform RUNSTATS and REBIND after, N— No
REBIND	Y/N To bind all related plans/packages

REPORT Utility

Overview

This unit introduces the REPORT utility. It covers availability of the REPORT utility and contains information about the values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the REPORT utility

The REPORT utility runs DB2 reports. This utility produces two types of reports: current or summary recovery information for a table space or table space set.

Availability of the REPORT utility

The REPORT utility is available from the Table Space panel.

Authority required to use the REPORT utility

If you want to use the option for the REPORT utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
 - DBADM or DBCTRL authority for the database
 - RECOVERDB privilege for the database
-

Values to provide that are used by multiple utilities

To use REPORT, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart (if desired)	“Values for Identifying and Restarting a Utility” on page 431

Specific values to provide when you use the utility

When you use the REPORT utility, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
RECOVERY	Y for REPORT recovery N for Tablespace set

Additional values required

Condition	Action to Perform
<p>If you specified Y for REPORT Recovery</p> <p>The values you specify for the Recovery? field determines additional values required. (The values can be specified in the profile.)</p>	<p>Type a value in as many fields as you want:</p> <p>CURRENT? Y for the latest SYSCOPY/ SYSLGRNG/ BSDS only.</p> <p>SUMMARY? Y for the latest SYSCOPY/ ARCHIVES volsers only.</p> <p>LOCALSITE Y for a report of LOCAL site records only.</p> <p>RECOVERYSITE Y for a report of RECOVERY site records only.</p>

RUNSTATS Utility

Overview

This unit introduces the RUNSTATS utility. The unit covers availability of the RUNSTATS utility. It also contains a chart listing the names of fields and the values you provide when you use the utility with either table spaces or index spaces. It also contains values specific to table spaces and additional instructions. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the RUNSTATS utility

The RUNSTATS utility gathers information about table and index space utilization and the efficiency of indexes. The system stores this information in the DB2 catalog.

Availability of the RUNSTATS utility

The RUNSTATS utility is available from these object list panels:

- Table Spaces panel
 - Indexes panel
-

Authority required for the RUNSTATS utility for table spaces

If you want to use the option for the RUNSTATS utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
- DBADM, DBCTRL, or DBMAINT authority for the database
- STATS privilege for the database

Values to provide that are used by multiple utilities

To use RUNSTATS, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for DB2 commands (Values for the CMDS BEFORE or CMDS AFTER fields)	“Options for DB2 Commands with Utilities” on page 426
Options for job control	“General Options for Job Control” on page 428
Values to identify the utility and restart (if desired)	“Values for Identifying and Restarting a Utility” on page 431

Specific values to provide when you use the utility with table spaces or indexes

When you use the RUNSTATS utility with table spaces or indexes, you provide the values for the fields shown in the chart. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
SHRLEVEL	R — REFERENCE (allow only read processing) or C — CHANGE (allow updates while utility runs)

Additional values to provide with table spaces only

When you use the RUNSTATS utility with table spaces, you also provide these values.

Field Name	Description on the Panel
INDEX	* for all indexes blank for none
REPORT	Y/N produce report or not
UPDATE	Update catalog A — All P — Accesspath S — Space N — None
TABLE STATS?	Y — Gather statistics for selected columns N — Do not gather table statistics * — Gather statistics on all columns and all tables
SAMPLE	The percentage of rows to examine when collecting non-indexed column statistics Note: Only available in DB2 Version 5.1 or higher
KEYCARD	Y — Collect key cardinalities on the concatenation of key columns for the specified indexes N — Do not collect key cardinalities on the concatenation of key columns for the specified indexes Note: Only available in DB2 Version 5.1 or higher

Additional values to provide with table spaces only (continued)

This chart continues with values you provide when you use the RUNSTATS utility with table spaces.

Field Name	Description on the Panel
NUMCOLS	The number of key columns to concatenate when collecting frequent values from the specified index Note: Only available in DB2 Version 5.1 or higher
COUNT	The number of frequent values to collect Note: Only available in DB2 Version 5.1 or higher

STOSPACE Utility

Overview

This unit introduces the STOSPACE utility. It covers availability of the STOSPACE utility and contains information about the values you provide when you use the utility. (The unit “Overview of the Chapter” on page 364 is a prerequisite for this unit.)

Introduction to the STOSPACE utility

The STOSPACE utility accumulates the total DASD allocated to each storage group in the DB2 catalog.

Availability of the STOSPACE utility

The STOSPACE utility is available from the Storage Groups panel.

Authority required to use the STOSPACE utility

If you want to use the option for the STOSPACE utility, you must have one of these authorizations:

- SYSADM or SYSCTRL authority
 - STOSPACE privilege
-

Values to provide that are used by multiple utilities

To use STOSPACE, you also provide these values used by multiple utilities. If you are unfamiliar with these types of values, you can find information in the source shown in the chart.

Types of Values	Source to Use
Options for job control	“General Options for Job Control” on page 428

Specific values to provide when you use the utility

There are no values specific to the STOSPACE utility. When you use the utility, the utility displays the Restart Options panel.

Chapter 17. Panels and Fields Used with Multiple Utilities

Introduction

This chapter reviews the panels and fields used with multiple utilities. After you have identified the values specific to the utility you want to use, then you can use the chapter to identify those fields that are used by the utility you want and by other utilities as well.

Chapter contents

Overview of the Chapter	420
Online Help and This Chapter	423
Partitioned Objects and Multiple Objects and Utilities	424
Options for DB2 Commands with Utilities	426
General Options for Job Control	428
Specifying Values for the Job Card and Job Statements	429
Values for Identifying and Restarting a Utility	431
Values for the Work Data Sets	432
Values for Sort Work Data Sets	435
Values for Copy Data Sets (LOCAL)	438
Values for Copy Data Sets (RECOVERY)	440
Values for Partitions	442
Modifying the Values for Device (Esoteric) Names	445
How !DB/WORKBENCH Determines Data Set Sizes	446

Overview of the Chapter

Background about the fields and panels and the utility profiles

You can begin to use DB2 utilities and other options without defining values for the fields used with the utilities. !DB/WORKBENCH provides a default utility profile that it uses if you do not create your own.

Prerequisites for using this chapter

Before using any of the utilities in this chapter, review the information in “Introduction to DB2 Utilities and Options” on page 339 and in “DB2 Utilities and Options Available with !DB/WORKBENCH” on page 363.

Organization of information in the chapter

This chapter provides information about the panels and fields that are used with many of the utilities. The units of information are covered in this order: first, fields that are included on panels specific to utilities, and then panels in the order provided on the menu listing options for profile utilities. Each unit of information covers these topics: the names of utilities that use the fields or panels, the connection between the utility profile options and the fields on panels, and the values for the fields themselves.

Relationship of this chapter to other chapters in this section

This chapter is intended to be used in conjunction with “DB2 Utilities and Options Available with !DB/WORKBENCH.” The prior chapter covers information specific to each of the utilities and options available with !DB/WORKBENCH. This chapter covers information common to most of the utilities. It also contains information about the values to use when you access a utility using the UTIL command and use the utility with partitioned objects, multiple objects, or both partitioned objects and multiple objects.

For example, you can review prerequisites for using the CHECK DATA utility in the unit “CHECK DATA Utility” on page 367. Depending on your experience with other !DB/WORKBENCH utilities or your satisfaction with your current utility profile, you can decide whether you also want to review prerequisites covered in “Options for DB2 Commands with Utilities” on page 426, “General Options for Job Control” on page 428, “Values for Identifying and Restarting a Utility” on page 431, and “Values for the Work Data Sets” on page 432.

Organization of this chapter and your needs

Review the chart for the information you want for the utility you plan to use. (If you are using partitioned data, also use “Values for Partitions” on page 442.)

	DB2 Commands	Partitioned and Multiple Objects	General Options	Restart Options	Work Datasets	Sort Work Datasets	Copy (LOCAL)	Copy (RECOVERY)	Partitions	Device Names	Dataset Sizes
	p.426	p.424	p.428	p.431	p.432	p.435	p.438	p.440	p.442	p.445	p.446
CHECK DATA	√	√	√	√	√	√				√	√
CHECK INDEX	√	√	√	√	√	√				√	√
COPY	√	√	√	√			√	√	√		
DIAG ABENDS											
DIAG SYSUTIL			√								
DISASTER REC			√				√	√			
ICCAT			√								
MERGE COPY	√	√	√	√	√		√	√	√	√	√
MODIFY			√	√					√		
QUIESCE	√		√	√							
REBIND			√								
REBIND PKGS			√								
REBIND PLANS			√								
RECOVER	√	√	√	√	√	√			√	√	√
REORG	√	√	√	√	√	√			√	√	√
REPORT			√	√					√		
RUNSTATS	√		√	√							
STOSPACE			√	√							

Additional resources for using utilities efficiently

You can create and set utility profiles so that you do not have to provide the values for the utility each time you run the utility. For more information, see “Creating or Modifying Profiles for DB2 Utilities and Options” on page 449.

You can also create user variables to include in your utility profiles. For more information, see “Using User Variables in Utility Profiles” on page 463.

Online Help and This Chapter

Reminder about resources for completing values

If you want to know more information about a specific value on a panel, type **Help** on the command line and press Enter or press the appropriate function key. The online Help panel is typical of the resources available.

```

Cmd ==>
----- HELP FOR UTILITY DATA SETS (1 of 2) -----
Press END to return, ENTER to continue
or TAB to highlighted text and press ENTER to display more information.

Description      Use this panel to specify the data sets to use for utilities.

DDname           Type the data definition name you want assigned in the the
                  JCL DD statement.

                  If you want to use the name from the skeleton JCL, type
                  SYSUT1, SYSREC, or SORTOUT.

                  If you want to use the table space or index space name as the
                  ddname, leave the field blank.

                  Otherwise, type #. The utility generated ddnames in the form
                  WORK0001 or UNLD0001 or SOUT0001.

Warning          If you type # or leave the field blank and use the Utility
                  Skeleton JCL, you must modify the JCL repeating section
                  to use the correct ddnames.

-----
GLOSSARY  INDEX  COMMANDS  FILTERS  WINDOWS  SELECTS  MORE->

```

Partitioned Objects and Multiple Objects and Utilities

Overview

This unit contains information about the values you use on panels when you access specific utilities using the UTIL command and use the utility with partitioned objects, multiple objects, or both partitioned objects and multiple objects.

If you are using utilities with partitioned objects and multiple objects

This unit is required if you are using the UTIL command and these utilities with partitioned objects, multiple objects, or both partitioned objects and multiple objects:

- CHECK DATA
- CHECK INDEX
- COPY
- MERGE COPY
- REORG
- RECOVER

Background about partitioned objects and multiple objects and job control

When you use the CHECK DATA, CHECK INDEX, COPY, MERGE COPY, REORG, and RECOVER utility with partitioned objects and multiple objects, you provide values on panels that are used by multiple utilities. One of the values you provide is the value for the type of job you want to use for the utility. The type of job you want to use determines the values you use on the other panels required for the utility.

Values to use for the utilities with partitioned objects and multiple objects

Review the chart to determine the values to use on panels when you use the UTIL command and the CHECK DATA, CHECK INDEX, COPY, MERGE COPY, REORG, or RECOVER utility with partitioned objects and multiple objects.

Task You Want to Perform	DDname Fields on Work Data Sets Panel	DDname Field on Copy Data Sets Panel	DSNUM/PART Field on Partition Specifications Panel	Job Control Field on General Options Panel
Use one job for all partitions	#	#	All	1
Use one job with one step for each partition	#, blank, or default	#, blank, or default	Each or some	2
Use one job for each partition	#, blank, or default	#, blank, or default	Each or some	3
Use one job for all objects	# or blank	# or blank	Not applicable	1
Use one job with one step for each object	#, blank, or default	#, blank, or default	Not applicable	2
Use one job for each object	#, blank, or default	#, blank, or default	Not applicable	3
Use one job for all partitions and all objects	#	#	All	1
Use one job with one step for each partition and each object	#, blank, or default	#, blank, or default	All, each, or some	2
Use one job for each partition and each object	#, blank, or default	#, blank, or default	All, each, or some	3

Options for DB2 Commands with Utilities

Utilities and options for DB2 commands

If you use these DB2 utilities, you provide the values for fields on the panel for options for DB2 commands: CHECK DATA, CHECK INDEX, COPY, MERGE COPY, QUIESCE, RECOVER, REORG, and RUNSTATS utilities.

Utility profiles and the DB2 commands

Unlike the values for other options in this chapter, the values are not shown on a separate panel when you create utility profiles. The options for the DB2 commands appear on the specific panel for the utility.

Values for CMDS BEFORE

The CMDS BEFORE option determines the DB2 command to be generated in a first step that precedes all the steps for a utility. Review the chart for acceptable values for the options.

Value Possible for the Field	Consequence of Using the Value
P	Generate STOP DATABASE commands for each data space to be processed (Caution: Most DB2 utilities require that the table space or index space not be STOPped.)
R	Generate START DATABASE ACCESS(RO) (Read Only) commands for each data space to be processed
U	Generate START DATABASE ACCESS(UT) (Utility) commands for each data space to be processed (UT allows only DB2 utilities to process against the data spaces.)
V	Combines P and U. Generate STOP DATABASE commands first and then START DATABASE ACCESS(UT) commands.

Values for CMDS AFTER

The CMDS AFTER option determines the DB2 command to be generated in a last step that follows all the utility step or steps. Review the chart for acceptable values for the options.

Value Possible for the Field	Consequence of Using the Value
P	Generate STOP DATABASE commands for each data space to be processed
R	Generate START DATABASE ACCESS(RO) (Read Only) commands for each data space to be processed
W	Generate START DATABASE ACCESS(RW) (Read and Write) commands for each data space to be processed
Z	Combines P and W. Generate STOP DATABASE commands first and then START DATABASE ACCESS(RW) commands

General Options for Job Control

Utilities and options using the values

If you use these DB2 utilities, you provide the values for fields shown on the General Options panel: CHECK DATA, CHECK INDEX, COPY, DIAG SYSUTIL, ICCAT, MERGE COPY, MODIFY, QUIESCE, RECOVER, REORG, REPORT, RUNSTATS, and STOSPACE utilities. REBIND, REBIND PKGS, and REBIND PLANS also use these fields.

Prerequisite for partitioned objects and multiple objects and utilities

If you are using the CHECK DATA, CHECK INDEX, COPY, MERGE COPY, RECOVER, or REORG utility with partitioned objects or multiple objects, the unit “Partitioned Objects and Multiple Objects and Utilities” on page 424 is a prerequisite for this unit.

Values on the General Options panel

These values are listed on the General Options panel. (The fields are listed in the order they appear on the panel.)

Field Name	Description on the Panel
DB2 SYSTEM	DB2 Subsystem identifier
DB2 STEPLIB	Use quotes if not prefixed with your userid
MODEL JCL	Utility JCL skeleton member name
REGION	Region size for EXEC (including K or M)
JOB CONTROL	1 - one step in one job for all objects 2 - multiple steps in one job for all objects (maximum of 255 steps per job) 3 - separate job for each object
MIGRATED	S to skip I to include migrated spaces
SYSOUT CLASS	Class for JES output (alphanumeric or *)
JOB CARD	User skeleton member name for Job Card
JOB STATEMENTS	Job name and related information

Specifying Values for the Job Card and Job Statements

Overview

This unit describes some considerations for specifying values for the job card and job statements.

Background about the value for the Job Card and the Job Statements

If you specify a value for the Job Card field, you do not have to provide entries in the Job Statements field. If you specify a value in the Job Card field and you provide entries in the Job Statements field, !DB/WORKBENCH requires you to make entries in the first four lines of the Job Statements field. It does not validate the entries.

Background about the Job Statements entry field

Use the optional Job Statements field to enter valid job card information (job name and related information) for the utility you are using. The first line of the four line Job Statements field controls the job name increment number if you want to generate a separate job name for each job generated.

Incrementing the job name

In the Job Control field, if you choose option 3, !DB/WORKBENCH creates a separate job for each object. If you choose option 3 and you specify a job name in the Job Statements field, !DB/WORKBENCH will increment each job name in the jobstream.

Prerequisites for incrementing a job name

If you want to increment the job name as described in this unit, you must observe the following rules:

- The “userid x ” must match the currently logged on USERID.
- The four lines of the Job Statements field must not change during the batch or online session. Thus, you should *not* use these !DB/WORKBENCH processing options during the session:
 - Option 15 (PROFILES) on the Utilities Menu
 - Changing the profile using the Set Profile Name field
 - During batch operations, by supplying the statement UTIL OPTS *profile_name*.

Generating unique names for multiple jobs for a single activity

Use this procedure to generate a unique name for each job when running multiple jobs for a single activity; for example, when running multiple image copies.

Incrementing the job name in online and batch processing is controlled by the first line of the four line Job Statements field on the General Utility Options (KTWPUTSV) panel. To control this job name increment, you must supply the first line in the Job Statements field as “//userid x JOB...”. !DB/WORKBENCH extracts the character “ x ” from the “userid x ” and increments the job name. The first job name is “userid x ”, the second job name is “userid $x+1$ ”, the third is named “userid $x+2$ ”, and this logic continues throughout the job stream. The job name is incremented numerically (1–9) or alphanumerically (A–Z) according to the starting value of “ x ”. This method provides a maximum of 35 unique names. When this limit is reached, the counter restarts with 1.

!DB/WORKBENCH places the value of “ x ” in an ISPF variable “&J” to be used in a custom job card member. The four lines of the Job Statements field are saved from online session to online session, but they are never recalled from the ISPF user profile.

Generating unique names for multiple jobs for multiple activities

To generate unique names when creating multiple jobs for multiple activities, for example, a REORG followed by an image copy for a group of table spaces, take these actions:

- Imbed ISPF variables such as &TSNAME or &J in a skeleton.
- Include the &J variable in the name of the skeleton.
- Specify the name of this skeleton in the JOB CARD field on the General Options panel.

This causes !DB/WORKBENCH to substitute incremental values for &J in the skeleton name. The initial value for &J and the incrementing process occur as described in the previous section on generating unique names for multiple jobs for a single activity. You can create an unlimited number of job control cards that reference a skeleton for an ISPF variable.

Values for Identifying and Restarting a Utility

Utilities using the values

If you use these DB2 utilities, you provide the values for fields shown on the Utility Restart Options panel: CHECK DATA, CHECK INDEX, COPY, MERGE COPY, MODIFY, QUIESCE, RECOVER, REORG, REPORT, RUNSTATS, and STOSPACE utilities.

Utility profiles and these values

If you want to complete these values when creating utility profiles, select Restart Options on the menu listing options for utility profiles. (When you use utilities online, !DB/WORKBENCH displays these fields on the panels for the specific utility.)

Values on the Utility Restart Options panel

These values are listed on the Utility Restart Options panel. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
UTILITY-ID	Unique Utility ID, default is <i>userid.jobname</i>
START	Leave blank if not restarting this Utility. Else, restart at C Commit or P Phase.

Values for the Work Data Sets

Utilities using the values

If you use these DB2 utilities, you provide the values for fields shown on the Work Data Sets panel: CHECK DATA, CHECK INDEX, MERGE COPY, RECOVER, and REORG utilities.

Prerequisite for utilities for partitioned objects and multiple objects

If you are using the CHECK DATA, CHECK INDEX, COPY, MERGE COPY, RECOVER, or REORG utility with partitioned objects or multiple objects, the unit “Partitioned Objects and Multiple Objects and Utilities” on page 424 is a prerequisite for this unit.

Utility profiles and these values

If you want to complete these values when creating utility profiles, select Work Dataset Options on the menu listing options for utility profiles.

Background about device (esoteric) names

For a list of the device (esoteric) names !DB/WORKBENCH recognizes and tasks to perform to change a device name, see “Modifying the Values for Device (Esoteric) Names” on page 445 later in this chapter.

Background about utility sizing parameters

For information about the way in which dataset sizes are determined, see “How !DB/WORKBENCH Determines Data Set Sizes” on page 446 later in this chapter.

Values on the Work Data Sets panel

These values are listed on the Work Data Sets panel. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
DDname	# = WORK $nnnn$, blank = "spacem", default = SYSUT1
UNIT	Device type for WORK datasets
SIZE	Override WORK dataset allocation value
STORCLAS	No description SMS storage class for the utility datasets
DATACLAS	No description SMS data class for the utility datasets
MGMTCLAS	No description SMS management class for the utility datasets

Values on the Work Data Sets panel (continued)

The chart continues the values on the Work Data Sets panel.

Field Name	Description on the Panel
DDname	# = UNDLnnnn, blank = "spacenm", default = SYSREC
UNIT	Device type for UNLOAD datasets
SIZE	Override UNLOAD dataset allocation value size
STORCLAS	No description SMS storage class for the utility datasets
DATACLAS	No description SMS data class for the utility datasets
MGMTCLAS	No description SMS management class for the utility datasets
DDname	# = SOUTLnnnn, blank = "spacenm", default = SORTOUT
UNIT	Device type for SORTOUT datasets
SIZE	Override SORTOUT dataset allocation value
STORCLAS	No description SMS storage class for the utility datasets
DATACLAS	No description SMS data class for the utility datasets
MGMTCLAS	No description SMS management class for the utility datasets

Values for Sort Work Data Sets

Overview

This unit describes the process !DB/WORKBENCH uses to create sort files for use by the CHECK, RECOVER, and REORG utilities.

Values for the sort work datasets

When you use the sort work datasets, you provide the values shown in the chart.

Field Name	Description on the Panel
SORTDEVT	The devicetype used for sort work datasets. Ignored if sort work datasets are not dynamically allocated.
SORTNUM	Number of sort datasets to dynamically allocate

Background about sort files and the CHECK, RECOVER, and REORG utilities

To initiate the CHECK, RECOVER, and REORG utilities, !DB/WORKBENCH creates datasets for a SORTIN and SORTOUT file. The system also creates temporary datasets for SORTWK nn files. This chart explains how !DB/WORKBENCH generates temporary sort files when you define a utility profile or use the online facility to override the definition for the utility profile.

Background about sort files and the CHECK, RECOVER, and REORG utilities (continued)

IF you initiate the...	AND you type...	THEN...
<ul style="list-style-type: none"> ● CHECK INDEX utility for index spaces ● CHECK DATA utility for table spaces ● RECOVER INDEX utility for index spaces 	a value in both the SORTDEVT and SORTNUM fields	!DB/WORKBENCH passes the SORTDEVT and SORTNUM values to a sort routine. Result: The sort routine uses these values to allocate the SORTWK nn datasets dynamically.
<ul style="list-style-type: none"> ● REORG INDEX utility for index spaces 	nulls or blanks in the SORTDEVT field and a value in the SORTNUM field	!DB/WORKBENCH will require a value for the SORTDEVT field.
<ul style="list-style-type: none"> ● REORG utility for table spaces 	a value in the SORTDEVT field and either nulls or blanks in the SORTNUM field	!DB/WORKBENCH passes no values to the sort routine. Result: The sort routine uses the install defaults to allocate the SORTWK nn datasets dynamically.
	nulls or blanks in the SORTDEVT and SORTNUM fields	!DB/WORKBENCH dynamically generates two temporary datasets for sort work datasets. These datasets are twice the size of the permanent dataset for the SYSUT1 file.

Examples of the allocation of SORTWK datasets

This chart provides some examples of the way in which an example sort routine, DFSORT, allocates the work datasets.

Example Input Values	DFSORT Work File Allocation
Example for SORTDEVT and SORTNUM values: Type SYSDA in the SORTDEVT field and 8 in the SORTNUM field	DFSORT uses the values you entered to dynamically allocate eight SORTWK nn datasets.
Example for a SORTDEVT value: Type SYSDA in the SORTDEVT field and nulls or blanks in the SORTNUM field	DFSORT uses the install defaults to allocate the SORTWK nn datasets dynamically.
Example for unspecified values: Type nulls or blanks in the SORTDEVT and SORTNUM fields.	!DB/WORKBENCH dynamically generates two temporary datasets that are twice the size of the SYSUT1 file.

Values for Copy Data Sets (LOCAL)

Utilities using the values

If you use the COPY and MERGE COPY DB2 utilities, you provide the values for fields shown on the Copy Data Sets (LOCAL) panel.

Prerequisite for utilities for partitioned objects and multiple objects

If you are using the CHECK DATA, CHECK INDEX, COPY, MERGE COPY, RECOVER, or REORG utility with partitioned objects or multiple objects, the unit “Partitioned Objects and Multiple Objects and Utilities” on page 424 is a prerequisite for this unit.

Utility profiles and these values

If you want to complete these values when creating utility profiles, select COPY Dataset Opts. 1/2 on the menu listing options for utility profiles.

Background about device (esoteric) names

For a list of the device (esoteric) names !DB/WORKBENCH recognizes and tasks to perform to change a device name, see “Modifying the Values for Device (Esoteric) Names” on page 445 later in this chapter.

Background about utility sizing parameters

For information about the way in which dataset sizes are determined, see “How !DB/WORKBENCH Determines Data Set Sizes” on page 446 later in this chapter.

Values on the first Copy Data Sets panel

These values are listed on the Copy Data Sets panel, which deals with COPY fields for the primary output image copy and backup copy. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
DDNAME1	# = COPYnnnn, blank = "spacem", default = SYSCOPY
UNIT1	Device type for Copy output datasets
SIZE1	Override Copy output dataset size (tracks)
STORCLAS	No description SMS storage class for the utility datasets
DATACLAS	No description SMS data class for the utility datasets
MGMTCLAS	No description SMS management class for the utility datasets
STACK	All datasets on one tape? Y/N or blank
DDNAME2	# = COPYnnnn+1, blank = no copy 2
UNIT2	Device type for Copy output datasets
STORCLAS	No description SMS storage class for the utility datasets
DATACLAS	No description SMS data class for the utility datasets
MGMTCLAS	No description SMS management class for the utility datasets
STACK	All datasets on one tape? Y/N or blank

Values for Copy Data Sets (RECOVERY)

Utilities using the values

If you use the DB2 COPY OR MERGE COPY utility, you provide the values for fields shown on the Copy Data Sets (RECOVERY) panel.

Prerequisite for utilities for partitioned objects and multiple objects

If you are using the CHECK DATA, CHECK INDEX, COPY, MERGE COPY, RECOVER, or REORG utility with partitioned objects or multiple objects, the unit “Partitioned Objects and Multiple Objects and Utilities” on page 424 is a prerequisite for this unit.

Utility profiles and these values

If you want to complete these values when creating utility profiles, select COPY Dataset Opts. 2/2 on the menu listing options for utility profiles.

Background about device (esoteric) names

For a list of the device (esoteric) names !DB/WORKBENCH recognizes and tasks to perform to change a device name, see “Modifying the Values for Device (Esoteric) Names” on page 445 later in this chapter.

Background about utility sizing parameters

For information about the way in which dataset sizes are determined, see “How !DB/WORKBENCH Determines Data Set Sizes” on page 446 later in this chapter.

!DB/WORKBENCH applies these guidelines when setting sizing parameters for datasets:

- Primary sizing matches the allocated size of the object
- Secondary sizing is 12.5% of the primary sizing
- Sort work allocations are 200% of the dataset size for the object

Values on the second Copy Data Sets panel

These values are listed on the Copy Data Sets panel, which deals with COPY fields for the temporary dataset for intermediate merged output. (The fields are listed in the order displayed on the panel.)

Field Name	Description on the Panel
DDNAME3	# = RECO n nnn, blank = no primary recovery copy
UNIT3	Device type for Copy output datasets
STORCLAS	No description SMS storage class for the utility datasets
DATACLAS	No description SMS data class for the utility datasets
MGMTCLAS	No description SMS management class for the utility datasets
STACK	All datasets on one tape? Y/N or blank
DDNAME4	# = RECO n nnn+1, blank = no backup recovery copy
UNIT4	Device type for Copy output datasets
STORCLAS	No description SMS storage class for the utility datasets
DATACLAS	No description SMS data class for the utility datasets
MGMTCLAS	No description SMS management class for the utility datasets
STACK	All datasets on one tape? Y/N or blank

Values for Partitions

Utilities using values

With values for partitions, the object you select (whether in batch or online) determines whether !DB/WORKBENCH displays a panel asking for entry of a value for the DSNUM/PART field. For example, if you have four table spaces displayed on a panel and if one of them is partitioned, !DB/WORKBENCH displays the Partition Specifications panel.

Prerequisite for utilities for partitioned objects and multiple objects

If you are using the CHECK DATA, CHECK INDEX, COPY, MERGE COPY, RECOVER, or REORG utility with partitioned objects or multiple objects, the unit “Partitioned Objects and Multiple Objects and Utilities” on page 424 is a prerequisite for this unit.

Utility profiles and these values

If the DB2 utility that you want to use supports partition values and if you plan to use utilities with the partitioned objects, you can select Partition Options on the menu listing options available for utility profiles.

Values and instructions on the panel

The illustration shows the values and instructions available on the Partition Specifications panel.

```

UTILITY PROFILES ----- DB/WORKBENCH -- DB2=DB2 -----
Cmd ==>                                                    99/11/16 08:55

          P A R T I T I O N   S P E C I F I C A T I O N S
DSNUM/PART ==> 1 ALL      ALL to process all partitions together, or
                  EACH to process each partition separately, or
                  SOME to process the partitions identified below, or 2
                  a number to process the corresponding partition.
                  (This specification is ignored for non-partitioned spaces.)

If SOME, use any non-blank character to select the partitions to be processed.

- 1   - 2   - 3   - 4   - 5   - 6   - 7   - 8   - 9   - 10
- 11  - 12  - 13  - 14  - 15  - 16  - 17  - 18  - 19  - 20
- 21  - 22  - 23  - 24  - 25  - 26  - 27  - 28  - 29  - 30 3
- 31  - 32  - 33  - 34  - 35  - 36  - 37  - 38  - 39  - 40
- 41  - 42  - 43  - 44  - 45  - 46  - 47  - 48  - 49  - 50
- 51  - 52  - 53  - 54  - 55  - 56  - 57  - 58  - 59  - 60
- 61  - 62  - 63  - 64

Press ENTER when you have completed your specifications.
    
```

- 1** Area for entry of the values of any partitioned object
(You can use ALL, EACH, or SOME.)
- 2** Instructions provided on the panel
- 3** Additional instructions if you use SOME as the value for
DSNUM/PART field

Examples of the consequences of the values

Review the chart for examples of the consequences that occur with the different values for the field.

Action You Perform	Utility Statements !DB/WORKBENCH Generates
ALL in the DSNUM/PART field	COPY TABLESPACE ABC.DEF DSNUM ALL ...
EACH in the DSNUM/PART field	COPY TABLESPACE ABC.DEF DSNUM 01 ... COPY TABLESPACE ABC.DEF DSNUM 02 ... COPY TABLESPACE ABC.DEF DSNUM 03 ...
SOME in the DSNUM/PART field and select partitions 1 and 3 on the display	COPY TABLESPACE ABC.DEF DSNUM 01 ... COPY TABLESPACE ABC.DEF DSNUM 03 ...

Modifying the Values for Device (Esoteric) Names

Utilities and options using the values

If you use these DB2 utilities, you should review the information in this unit: CHECK DATA, CHECK INDEX, MERGE COPY, RECOVER, and REORG.

Device (esoteric) names recognized by !DB/WORKBENCH

!DB/WORKBENCH recognizes the following device (esoteric) names:

- 3420
 - 3490
 - CART9
 - REEL
 - TAPE
-

Providing support for alternate device names

To provide support for a different device name, make the modifications to !DB/WORKBENCH shown in this chart.

Task	Action: Change the device name in...
Modify the appropriate panels.	Panels KTWPUTC2 and KTWPUTC3
Modify the appropriate skeleton.	Skeleton KTWUTDVC
Modify the appropriate message.	Message KTWU010, found in member KTWU10.

How !DB/WORKBENCH Determines Data Set Sizes

Overview

You have various options for allocating datasets for the utilities. This unit discusses how your definitions are used by !DB/WORKBENCH to size the various datasets it uses during utility generation.

Utilities and options using the values

If you use these DB2 utilities, you should review the information in this unit: CHECK DATA, CHECK INDEX, MERGE COPY, RECOVER, and REORG.

Options available for allocating datasets

You can now allocate datasets in cylinders, tracks, or records (blocks). Use the SPCCTL variable in the KTWUTJCL skeleton to control this option for all of the DB2 utilities. This chart describes the sizing variables in the KTWUTJCL skeleton and the changes you can make.

Variable	Function	Additional Information
SPCCTL	Controls the type of dataset allocation. The values are <ul style="list-style-type: none"> ● 1 to allocate cylinders ● 2 to allocate tracks ● 3 to allocate records (blocks) 	The default is 2, which allocates datasets in tracks.
BLKSIZ	Globally controls the block size used by the utilities.	The default value (24576) supports the calculation of accurate sizing for image copies. Changes to this value can cause errors in allocating space for image copies.
REGION	Sets the default size (4M) for the region.	Provide a value if there is none for this required variable.

Size allocations determined by dataset type

The amount of space that is allocated for a data set is determined by the dataset type. Different size allocations are provided to accommodate these types of datasets:

- SYSUT1(work)
- UNLOAD
- SORT
- COPY and RECOVERY

Background about utility sizing parameters

!DB/WORKBENCH applies these guidelines when setting sizing parameters for datasets:

- Primary sizing matches the allocated size of the object
- Secondary sizing is 12.5% of the primary sizing
- Sort work allocations are 200% of the dataset size for the object

Sizing for incremental copies of migrated objects

Refer to the chart for an explanation of how parameters for the utility profiles manage sizing for incremental copies of migrated objects.

Task You Want to Perform	Parameter to Use	Additional Information
Accurately size the backup dataset for an incremental image copy	Type M in the IMAGECOPY field on the COPY Utility Options panel.	!DB/WORKBENCH reads the DB2 space map pages to calculate the accurate sizing for the copy dataset.
Include a migrated object in the incremental image copy	Type I in the MIGRATED field on the the General Options panel.	!DB/WORKBENCH does not perform accurate sizing for migrated objects. The sizing defaults to the allocated size that is set by DFHSM.
Exclude a migrated object from the incremental image copy	Type S in the MIGRATED field on the the General Options panel.	!DB/WORKBENCH skips any migrated objects.
Selectively choose migrated objects for an incremental image copy	Leave the MIGRATED field on the the General Options panel blank .	!DB/WORKBENCH displays the Migrated Data Set panel for each migrated object so that you recall only those datasets you want.

Sizing considerations for index recovery when an object is missing

When !DB/WORKBENCH generates JCL for the RECOVER utility, it checks to determine the existence of physical objects. However, if an object is missing, !DB/WORKBENCH still generates the JCL.

If an object is missing when the recovery is for an index, !DB/WORKBENCH generates the JCL but needs information on the work dataset size. To obtain a size to work with, !DB/WORKBENCH follows the procedures in this chart.

Step	Action
1	<p>!DB/WORKBENCH first attempts to obtain a dataset size from one of the following:</p> <ul style="list-style-type: none"> ● Work Data Sets panel (KTWPUOG2), which can be updated during utility profile customization ● Recover Index Utility Options panel (KTWPUTRX). You can access this panel by <ul style="list-style-type: none"> – First, entering the U select or UTIL command on the Indexes panel (KTWPDSIX) – Then, selecting the option to recover an index from the pop-up menu that !DB/WORKBENCH displays
2	<p>If the work dataset size has not been entered on one of these panels, !DB/WORKBENCH goes to the JCL skeleton KTWUTDDW to retrieve the default value (150,18). It issues message KTWU067N:</p> <p style="text-align: center;">One or more index datasets missing. Default sizing will be used.</p>

Note: The size allocation depends on the SPCCTL variable in the KTWUTJCL skeleton, described earlier in this unit.

Chapter 18. Creating or Modifying Profiles for DB2 Utilities and Options

Introduction

This chapter introduces the utility profile as a method to provide the values required to run DB2 utilities and options without keying the values each time. It includes background and instructions for creating or modifying utility profiles. The chapter “Introduction to DB2 Utilities and Options” on page 339 is a prerequisite for this chapter.

Chapter contents

Overview of the Chapter	450
Background about the Profile Option on Utilities Menus	452
Accessing the Utility Profiles Menu	454
Beginning the Process of Creating a Profile	456
Background about Options and Panels in the Profile	458
Providing Values on Panels	460
Saving a Profile	462

Overview of the Chapter

Background about creating utility profiles

!DB/WORKBENCH provides a default utility profile that it uses if you do not create ones of your own. You can only create your own utilities profiles if you are authorized to update the system PDS.

Before you create your own profiles, you need to be familiar with the panels that the system displays when you use the utility and you need to understand how the utility works with your DB2 system. When you create a new utility profile, you may want to experiment with it online before applying it in batch processing.

Organization of information in this chapter

This chapter introduces the utility profile and provides instructions for using the options that create and set profiles on the Utilities Menu for object list panels.

Implications for batch processing

You also use utility profiles in batch processing. !DB/WORKBENCH uses one of these profiles when you perform batch processing:

- Your own utility profile that you have set (made active) on one of the utilities menus for an object list (such as on the Table Spaces Utilities menu)
- The default utility profile provided by !DB/WORKBENCH (if you have not made a profile and set it)
- The utility profile that you set in the batch program by including the command `UTIL OPTS profilename` in your job

Additional information about batch and utility profiles

If you want additional information about utility profiles and batch, see “JCL for Utilities, Batch Jobs, and Extracts” on page 521.

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Utilities Menus	Utility Profile	Beginning Creation	Options for Profiles	Providing Values	Saving Profiles
	p.452	p.454	p.456	p.458	p.460	p.462
Review options available for utility profiles	√					
Review methods for creating a utility profile (and the differences in the contents of the profile)			√			
Review the options and utility panels that you can include in the profile				√		
Create a new utility profile and make it the active profile (without modifying values in the profile)	√	√	√			√
Create a new utility profile, modify values in the profile, and make the profile active	√	√	√	√	√	√

Reminder about resources

For information about specific values to provide on panels required for DB2 utilities and options, see online Help for the panel. To review all prerequisite values before you begin, see “DB2 Utilities and Options Available with !DB/WORKBENCH” on page 363 and “Panels and Fields Used with Multiple Utilities” on page 419.

Background about the Profile Option on Utilities Menus

Overview

This unit provides background needed to use the Profile option on the Utilities Menus available from object list panels.

Considerations before creating a utility profile

Before you decide to create a utility profile, you should be familiar with the process of generating utility statements used by !DB/WORKBENCH. You should understand the relationship between the utility options and the utility skeletons. For information on the utility skeletons, see “JCL for Utilities, Batch Jobs, and Extracts” on page 521.

Location for the utility profile

!DB/WORKBENCH stores utility profiles in your user ISPF profiles. In this way, when you use a profile and then later return to the system, you can use the same profile as you did previously.

Availability of the default utility profile

!DB/WORKBENCH must have an active utility profile. However, you do not have to create a profile in order to use the utilities. !DB/WORKBENCH provides a default utility profile (named DUMMY) that includes defaults for most values for all of the supported utilities. If you do not specify a different utility profile, !DB/WORKBENCH uses the default profile.

Possibility of multiple profiles

You can only use one utility profile at a time; however, you can have multiple utility profiles saved in your user ISPF profile. For example, you can have multiple profiles for different purposes when using the utilities.

Utility profiles and options on the Utilities Menu

Each of the object list panels that supports utility functions also supports a Utilities Menu. For example, if you are on the Table Spaces panel and use the UTIL command, !DB/WORKBENCH displays the Table Spaces Utilities Menu. The menu lists utilities available with the object and provides two options on those menus that are significant in creating and setting (making active) utility profiles. (The numbers associated with the options vary with the object list.)

Option on the Utilities Menu	Description on the Panel	Purpose
PROFILE	Create Utility Profile	To display existing profiles for copying, selecting, or updating To begin the process of creating a new profile
SET PROFILE (and a field showing a name of a profile)	Set Utility Profile	To show the current, active profile To make another profile active

Reminder about accessing a utility profile

If you need a reminder about how to access one of the utility profiles or how to set a profile, see “Setting the Active Utility Profile” on page 351 in the chapter “Introduction to DB2 Utilities and Options.”

Accessing the Utility Profiles Menu

Overview

This unit includes instructions for accessing the Utility Profiles Menu. You use the unit whether you want to create a new profile or copy or modify an existing profile.

Procedure to follow to access the Utility Profiles Menu

Follow these steps to begin to create a utility profile.

Step	Action
1	On the Utilities Menu for an object list panel, type the number of the Profile option in the Option field.
2	<p>Press Enter.</p> <p>Result: The system displays the Utility Profiles Menu, which lists the names of the current utility profiles.</p> <div data-bbox="571 951 1250 1249" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <pre> UTILITY PROFILES ----- DB/WORKBENCH -- DB2=D31A----- PROFILE 1 OF 18 Cmd ==> Scroll ==> PAGE Below is a list of all defined Utility Profiles. Select (S) one in order to display or update the Profile. You may also erase entries with (E) if you were the creator. SEL NAME CREATOR LAST CHG --- - <-- ENTER NAME HERE TO CREATE A NEW UTILITY PROFILE - JTC240 TDD856 10/10/99 - JTC250 TDD856 10/28/99 - JTC260 TDD856 10/17/99 ***** BOTTOM OF LIST ***** </pre> </div>

Object list panels for utilities

To use the options available on the Utility Profiles Menu, you must begin the procedures on an object list panel that supports utilities. To determine the object list panels that are available from the !DB/WORKBENCH Primary Menu and that provide the UTIL command, see the chapter “Dictionary of Commands and Selects for Object List Panels” in *!DB/WORKBENCH Reference Manual*.

On the Tables panel, you can also use utilities on object list panels other than on those panels available from the !DB/WORKBENCH Primary Menu. For example, you can use utilities on the panel displayed by using the TSSET command.

Options available with the Utility Profiles Menu

Review the chart for the options appropriate for the task you want to perform.

Task You Want to Perform	Method	Name
Erase a profile	Select	E
Move the name of a profile to the top of the panel	Command	L <i>profilename</i>
Print the entire list to a JES printer	Command	PRNT
Select a utility profile to copy, display, or update	Select	S

Restrictions on using options available with the menu

No restrictions exist if you want to move the name of a profile or print the list of profiles. You also do not need authority to select a utility profile or copy it, or select a profile and display all of the panels associated with it.

These actions, however, do have restrictions. You must be the creator of the utility profile to perform these actions:

- Erase an existing utility profile
- Update an existing utility profile

Determining what to do next

If you want to create a new utility, follow the instructions in the unit “Accessing the Utility Profiles Menu” on page 454.

Beginning the Process of Creating a Profile

Overview

This unit provides instructions about two methods for creating a new utility profile and discusses the contents of the profiles depending upon the method you use.

Background about the name of a new utility profile

Whether you are creating a new profile panel-by-panel, or by copying an existing profile, you must provide a name that is no longer than six alphanumeric characters. The area for entry of the name is indicated by an arrow pointing to the field that is under the column entitled Name.

Background about the methods you use

Review the chart to determine the values that will be in the profile depending on the method you use.

Method you Use to Create the Profile	Values in the Utility Profile
Select an existing utility profile on the Utility Profiles Menu and perform additional steps.	The values in the utility profile you selected
Do not select an existing utility profile.	The values in whatever utility profile is currently set (active)

Creating a new profile by selecting from the menu

If you want to select an existing utility profile as the basis for your new profile, follow these steps.

Step	Action
1	Type S in the field for the profile that you want to use.
2	Press Enter. Result: The system displays the menu listing options for the profile.
3	Type END on the command line and press Enter, or press the appropriate function key. Result: The system redisplay the Utility Profiles Menu.
4	Type the new name of the profile in the field under the column entitled Name.
5	Press Enter.

Creating a new profile without selecting from the menu

If you want to use the values in the utility profile that is currently active (set) as the basis for your new profile, follow these steps.

Step	Action
1	Type the name of the new utility profile in the field under the column entitled Name.
2	Press Enter.

Determining what to do next

If you want background about how the options for utility profiles work together see the unit “Background about Options and Panels in the Profile” on page 458. If you want instructions for using the panels (including for checking of errors), see “Providing Values on Panels” on page 460.

Background about Options and Panels in the Profile

Overview

This unit reviews the options available for utility profiles and shows which utilities require profiles that are used by other utilities.

Overview of options available on the menu

During the process of creating a utility profile, !DB/WORKBENCH displays this menu listing options for utility profiles.

```

----- DB/WORKBENCH -- DB2=D31A-----
Cmd ==>

Utility Profile: NEWPRF
Press ENTER to display the options in order or select by number. To start
at a topic and then continue from there, select that entry by number,
press ENTER. To return to this panel, press ENTER again. To terminate
the generation or update of a Utility Profile, return to this menu by
pressing END, type CANCEL in the Cmd area and press ENTER.

      GENERAL                                UTILITY DEFAULTS
-----
1 - General Opts.                          8 - CHECK INDEX          18 - REPORT
2 - Restart Opts.                          9 - CHECK DATA         19 - DIAGNOSE
3 - Work Dataset Opts.                     10 - COPY                20 - INDEX RECOVER
1 4 - COPY Dataset Opts.                   11 - MERGECOPY           21 - INDEX REORG
5 - RECOVERY Dataset Opts.                 12 - MODIFY              22 - INDEX RUNSTATS
6 - Partition Opts.                        13 - QUIESCE             23 - PLAN REBIND Opts.
2 7 - User Vars.- Not Avail.               14 - TBLSPC RECOVER 1/2  24 - PKG. REBIND Opts.
                                           15 - TBLSPC RECOVER 2/2  25 - REBIND Owner Opts.
                                           16 - TBLSPC REORG
                                           17 - TBLSPC RUNSTATS
                                           |
                                           3

```

- 1** Panels listing values required for many of the utilities
- 2** Panel to be used for variables you want to include for your site (For instructions, see “Setting the Active Utility Profile” on page 351.)
- 3** Panels listing values required for specific utilities available

Using profile panels common to many utilities

Many utilities require the same values as other utilities. You can complete all of the panels, or you can complete only those panels required for specific utilities. The chart indicates which common profile panels are required for each utility and shows the name of the utilities menu where you can access the utility. You complete the Partition Options panel only if you have partitioned table spaces.

	General Options	Restart Options	Work Dataset Options	Copy Dataset Options 1/2	Copy Dataset Options 2/2	Partition Options	User Variables
CHECK DATA (Table Spaces)	√	√					√
CHECK INDEX (Table Spaces; Indexes)	√	√	√				√
COPY (Table Spaces)	√	√		√	√		√
DIAG ABENDS (Table Spaces; Indexes)							
DIAG SYSUTIL (Table Spaces)	√						
ICCAT (Table Spaces)	√						
MERGE COPY (Table Spaces)	√	√	√	√			√
MODIFY (Table Spaces)	√	√					
QUIESCE (Table Spaces)	√	√					√
REBIND (Table Spaces; Indexes)	√						
REBIND PKGS (Table Spaces; Indexes)	√						
REBIND PLANS (Table Spaces; Indexes)	√						
RECOVER (Table Spaces; Indexes)	√	√					√
REORG (Table Spaces; Indexes)	√	√	√				√
REPORT (Table Spaces)	√	√	√				
RUNSTATS (Table Spaces; Indexes)	√	√					√
STOSPACE (Stogroups)	√	√					

Reminder about utilities menus

If you need a reminder about the variety of methods for accessing utilities menus, see the unit “Accessing the Utility Profiles Menu” on page 454.

Providing Values on Panels

Overview

This unit covers how to provide values on panels and how to use !DB/WORKBENCH to check for errors in the values for your profiles. It also covers how you can cancel the creation of a profile if necessary.

Background about the Enter key and checking for errors

!DB/WORKBENCH can check the values you have provided in the utility profile to determine whether they are valid entries. Whether !DB/WORKBENCH checks for errors depends on the function you use after you complete the values for the panel.

Action you Want !DB/WORKBENCH to Perform	Your Action
Check for errors in the panel and display a message so you can correct errors	Press the Enter key.
Save the values (but do not check for errors)	Type END on the command line and press the Enter key, or press the appropriate function key.

Background about order of display of panels for the profile

When you select an option, the system displays the panel for the option you selected. When you complete the values on the panel and press Enter, the system displays the next panel in the order shown on the menu listing options for the profile.

Reminder about online Help and panels

Online Help provides information about the values to provide for each field on the panels displayed with the options.

Selecting and completing a panel for the utility profile

Follow these steps to select an option to include in the profile.

Step	Action
1	On the menu listing options for the profile, type the number of the option on the command line.
2	Press Enter. Result: The system displays the panel for the option you selected.
3	Complete the values on the panel.
4	Press Enter. Result: The system displays the next panel in the series.
5	Perform the appropriate action. <ul style="list-style-type: none"> ● If you want to complete another panel, press Enter. ● If you want to return to the menu listing options for the profile, press the Enter key twice.

If you decide not to complete creation of a profile

If you decide you do not want to complete the creation of a profile, follow these steps.

Step	Action
1	Type END on the command line and press the Enter key, or press the appropriate function key. Result: The system displays the menu listing options for the profile.
2	On the menu, type CANCEL on the command line.
3	Press Enter. Result: !DB/WORKBENCH does not save the profile. (The name you gave the profile is not listed on the Utility Profiles Menu.)

Saving a Profile

Overview

This unit provides instructions for saving a utility profile.

Background about saving the utility profile

For a utility profile to be saved, you must perform two actions:

- Save the profile using instructions in this unit.
 - Exit !DB/WORKBENCH successfully. (If you ABEND, the system does not save the profile.)
-

Procedure to save a utility profile

Follow these steps to save a utility profile after you have completed all values on panels for options you want.

Step	Action
1	<p>On the menu listing options for the profile, type END on the command line and press Enter, or press the appropriate function key.</p> <p>Result: The system displays the Utility Profiles Menu with the name of the new profile in its correct alphabetical listing among all the profiles.</p>
2	<p>Type END on the command line and press Enter, or press the appropriate function key.</p> <p>Result: The system redisplay the Utilities Menu for the object list panel where you began the process of creating a utility profile and displays the name of your profile in the Set Profile field.</p>

If you decide to resume creation of the profile

When you return to the menu listing options for the profile, you can save, cancel, or resume using the options. To resume, perform either of these actions:

- Type the number of an option on the command line and press the Enter key.
- Press the Enter key. (!DB/WORKBENCH displays the next panel in the order shown on the menu listing options for the profile.)

Chapter 19. Using User Variables in Utility Profiles

Introduction

This chapter provides background and prerequisites for creating user variables. It contains an example that shows how to create user variables. The chapter also contains information about using utilities and the panel for user variables.

Chapter contents

Organization of the Chapter	464
Introduction to Creating User Variables	466
Overview of the Process of Creating and Using User Variables	467
Prerequisites for Creating User Variables	468
Background about the Template for the Panel for User Variables	470
Background about Variables in the Template	472
Characteristics of the Example of Modifying the Template	474
Example of Modifying the Template	476
Introduction to Creating Help for User Variables	478
Prerequisites for Creating Help for User Variables	479
Background about Templates for Help for User Variables .	480
Characteristics of the Example of Creating Help for User Variables	482
Example of Creating Help for User Variables	484
Completing the Process of the Example	486

Organization of the Chapter

Background about user variables and utility profiles

User variables with utility profiles are optional. If you want variables in utility profiles that are distinctive to your site, you can create them by following the instructions in this chapter.

!DB/WORKBENCH provides an option to show any existing user variables when you or a user creates a utility profile. If variables are complete, !DB/WORKBENCH displays the panel for user variables along with other profile panels. (It is available from option 7 on the Utility Profile menu.)

Organization of this chapter

This chapter provides prerequisites and instructions for creating user variables. This chapter also provides prerequisites and instructions for creating online Help for the user variables you create. It uses an example throughout the chapter to show how you can create the user variables and the online Help.

Relationship between this chapter and other chapters

These chapters are prerequisites for this chapter: “Creating or Modifying Profiles for DB2 Utilities and Options” on page 449 and “Introduction to DB2 Utilities and Options” on page 339.

Reminder if you want information about using variables in profiles

This chapter covers how to create the user variables. If you need information about using one of the panels available with utility profiles, see “Introduction to DB2 Utilities and Options” on page 339.

What this chapter does not cover

This chapter does not include information on using ISPF to work with panels or variables. For additional information on these topics, see the IBM documentation for ISPF.

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Introduction to User Variables	Process	Prerequisites for User Variables	Template for User Variables	Example of User Variables	Introduction to Online Help	Prerequisites for Online Help	Template for Online Help	Example of Online Help
	p.466	p.467	p.468	p.470	p.474	p.478	p.479	p.480	p.484
Review background about the panel for user variables	√	√							
Review prerequisites for creating user variables			√						
Review an example of creating user variables	√	√	√	√	√				
Review background about the online Help for user variables	√					√			
Review prerequisites for creating online Help for user variables	√						√		
Review an example of creating the online Help for the user variables you create						√	√	√	√

Introduction to Creating User Variables

Overview

This unit contains information about the templates and resources for creating user variables that !DB/WORKBENCH provides for the panel for user variables.

Background about the template provided within !DB/WORKBENCH

!DB/WORKBENCH provides the template for you to use if you want to create your own variables. The template panel includes ISPF values for these types of items:

- Color
 - The variables themselves
-

ISPF and creating the user variables

You can use all facilities available with ISPF when creating panels. For example, you can create variables with these features:

- Checking and validation of values
 - Default values
 - Cursor control
-

Panels with user variables and processing

Users can use the variables specific to your site either in batch or online. If a user completes the panel for user variables, !DB/WORKBENCH displays it when they create utility profiles and they can change the values just as they do for any profile values.

Overview of the Process of Creating and Using User Variables

Overview

This unit contains an overview of the process you follow when you create user variables. It also contains the process to follow to use the panel when a user uses the utilities.

Process of creating user variables

If you decide to create user variables for your utility profiles, you perform these actions.

Step	Action
1	Modify the template panel to create the user variables you want for your site. Example: You define a variable entitled <i>userdept</i> .
2	Modify the appropriate skeleton(s) to use the user variable you have created.

How the user can use the panel for user variables

If you create the user variables, the user of the utility can perform these steps.

Step	Action
1	The user selects a utility to use with an object. Result: !DB/WORKBENCH displays the panels with values needed for the utility, including the panel where the user provides values.
2	The user provides a value for the user variables you created. Example: The user types PRODUCTION in the user defined field <i>userdept</i> and presses Enter.
3	The user completes any other required values for the utility. Result: The system displays the Process Member menu.
4	On the Process Member menu, the user submits the utility. Result: !DB/WORKBENCH substitutes the values the user entered in all locations in the jobstream where you have used the variable.

Prerequisites for Creating User Variables

Overview

This unit contains prerequisites for when you create user variables. It also contains information about the utilities that support user variables.

Knowledge required when creating user variables

To create a user variable, you must have experience with ISPF panels and user variables. You also must be familiar with the ISPF editor. If you need information, see the IBM documentation on ISPF.

Authority required to create user variables

You must have these authorities to create user variables:

- Authority to update the panel library for !DB/WORKBENCH
 - Authority specific to your site (such as RACF or ACF2)
-

Utilities that can access user variables

You can create user variables for use with these utilities:

- CHECK DATA
 - CHECK INDEX
 - COPY
 - MERGE COPY
 - QUIESCE
 - RECOVER
 - REORG
 - RUNSTATS
-

Requirements for skeletons and user variables

You must add the variable to the skeleton where you want to use the user variable. For additional information, see the comments and instructions in the skeletons themselves and the chapter “JCL for Utilities, Batch Jobs, and Extracts” on page 521.

Panel library for the panel for user variables

The template for the panel for user variables is in the member KTWPUTUS in the panel library for !DB/WORKBENCH.

Protecting modifications to the panel library and skeletons

Be sure to protect your modifications to the panel library and skeletons. When you install subsequent versions of !DB/WORKBENCH, you overwrite your customized versions of the panels and the skeletons unless you protect these modifications. For example, you can copy !DB/WORKBENCH skeletons and make modifications to the copies.

Background about the Template for the Panel for User Variables

Overview

This unit reviews the elements of the template that you use if you want to create user variables for use within the utilities.

Elements of the template

The illustration shows the elements in the template that you use when you create your user variables.

```

)ATTR                                /* USER OPTIONS */
  TYPE(INPUT) CAPS(ON)                1
  @ TYPE(OUTPUT) INTENS(HIGH) CAPS(OFF) JUST(LEFT)
  .
  .
)BODY EXPAND(\\)
@&SFXSCRNM&SFXTITLE\-\
%Cmd%====>_ZCMD
+
+   Data Set (44) ^====>_LRGDSN          +
+   Option Info. ^====>_DULLDAY          +
+
?   ? To stop the display of this panel modify member KTWPUTUS in the product
?   ? panel library. Reference the .RESP line in the )INIT section.
+   + THIS IS A SAMPLE PANEL FOR DEFINING UTILITY VARIABLES UNIQUE TO YOUR 2
+   + INSTALLATION.

YOU DEFINE YOUR OWN VARIABLE NAME AS YOU WOULD FOR ANY OTHER ISPF PANEL.
YOU CAN SET A DEFAULT OR INITIAL VALUE FOR THIS VARIABLE IN THE )INIT
SECTION. THE VALUE MUST BE A 1-32 CHARACTER STRING.
.
.
)INIT ] 3
.CURSOR = LRGDSN 4
.HELP   = KTWHTUS
.
.
IF (&UTUSRMSG = &Z) 5
/* UNCOMMENT THE FOLLOWING LINE TO HAVE THE USER VARS. PANEL DISPLAY */
/* &UTUSRACT = YES */
.
.
&UTUSER01 = 'LRGDSN ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */ 6
&UTUSER02 = 'DULLDAY ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */
/* REMOVE COMMENTS TO SET INIT VALUE */
/*
/* IF (&LRGDSN = ' ',&Z) 7
/*   &LRGDSN = 'TDDB.DBTOOLS.DB2.DEMO.DSWMQ.WKB.SYSTEM' */
)PROC
VPUT (LRGDSN ) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */ 8
VPUT (DULLDAY) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */
)END

```

Elements of the template (continued)

The list defines the elements of the template.

- 1** Controls the characteristics of the panel
- 2** Provides initial instructions and samples of two user variables; serves as a location for your user variables and any instructions you want to provide (You can use line 16 to line 38 in the ISPF editor for your text and variables.)
- 3** Marks the end of the text that the system displays on the panel and the beginning of instructions for processing the data (To execute the panel successfully, be sure that the word *)INIT* is on line 39 or above.)
- 4** Controls the position of the cursor and the link to the Help panel
- 5** Controls whether the system displays the product panel for user variables
- 6** Provides samples of the definitions required to assign the name of the sample to one of the 32 profile variables *!DB/WORKBENCH* supports; serves as the location for your definitions for your user variables
- 7** Provides sample values for initialization for the samples of the user variables; serves as the location for the initialization instructions for your user variables
- 8** Directs the system to use the shared pool (*!DB/WORKBENCH* requires use of the shared profile pool; do not change the value for the pool.)

Background about Variables in the Template

Overview

This unit provides background about user variables in the template and requirements for values for user variables.

References to the user variables in the template

The illustration shows the references to the user variables in the template you use to create your user variables.

```

)ATTR                                /* USER OPTIONS */
  _ TYPE(INPUT) CAPS(ON)
  @ TYPE(OUTPUT) INTENS(HIGH) CAPS(OFF) JUST(LEFT)
  .
  .
)BODY EXPAND(\\)
@&SFXSCRNM&SFXTITLE\-\
%Cmd%====>_ZCMD
+                                     5
+   Data Set (44)  ^====>_LRGDSN      +
+   Option Info.  ^====>_DULLDAY      +

? To stop the display of this panel modify member KTWPUTUS in the product
? panel library. Reference the .RESP line in the )INIT section.
+
  THIS IS A SAMPLE PANEL FOR DEFINING UTILITY VARIABLES UNIQUE TO YOUR
  INSTALLATION.

  YOU DEFINE YOUR OWN VARIABLE NAME AS YOU WOULD FOR ANY OTHER ISPF PANEL.
  YOU CAN SET A DEFAULT OR INITIAL VALUE FOR THIS VARIABLE IN THE )INIT
  SECTION. THE VALUE MUST BE A 1-32 CHARACTER STRING.
  .
  .
)INIT                                6
.CURSOR = LRGDSN
.HELP   = KTWHUTUS
.
.
IF (&UTUSRMSG = &Z)
/* UNCOMMENT THE FOLLOWING LINE TO HAVE THE USER VARS. PANEL DISPLAY */
/* &UTUSRACT = YES */
.
.
7
&UTUSER01 = 'LRGDSN ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */
&UTUSER02 = 'DULLDAY ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */
/* REMOVE COMMENTS TO SET INIT VALUE */
/*                                     */
/* IF (&LRGDSN = ' ',&Z) */
/* &LRGDSN = 'TDDB.DBTOOLS.DB2.DEMO.DSWMQ.WKB.SYSTEM' */
)PROC 9
VPUT (LRGDSN ) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */
VPUT (DULLDAY) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */
)END

```

References to the user variables in the template (continued)

The list defines the references to the user variables in the template.

- 1** Text that !DB/WORKBENCH displays on the panel
- 2** Symbol used to indicate a field for entry (optional)
- 3** Indicator of the beginning of the field for entry by the user
- 4** Name for the user variable you are defining (maximum length is 8 characters; minimum length is 1 character)
- 5** Indicator for the end of the field for entry by the user (maximum length of the field is 32 characters)
- 6** Variable you specify to display the cursor
- 7** Assignment of the specified variable to one of the 32 profile variables supported by !DB/WORKBENCH
- 8** Instructions for initialization of the specified variable (You can use any methods available with ISPF. For example, you can specify a default value for the system to display or you can use ISPF user exits.)
- 9** Assignment of the specified variable to the shared profile pool

Requirements for values for user variables

The maximum length of a user variable is 32 characters. How you set up the variable determines the action !DB/WORKBENCH performs:

- If you set up the length of the variable to be 32 characters or less, !DB/WORKBENCH stops the user from entering more than 32 characters.
- If you set up the length to be 33 characters or more, !DB/WORKBENCH discards the user entry and replaces the entry with the default value or the value specified in the profile.

Characteristics of the Example of Modifying the Template

Overview

This unit gives characteristics of the example of modifying the template you use to create a panel for user variables. It shows the appearance of the panel and provides the values to use for the example. (For the example of modifying the template, see “Example of Modifying the Template” on page 476.)

Overview of the tasks in the example for creating user variables

The template you use to create the panel contains samples of two user variables. You modify the samples of user variables in the template. You modify the samples and create a third user variable in the template to create the panel. The panel for the example shows the three user variables.

Appearance of the panel for the example

The illustration shows the appearance of the panel using the example used in this chapter.

```
KTWPUTUS -----  
Cmd ==>  
  
Data Set Name ==> TDDB.DBTOOLS.DB2.DEMO.DSWMQ.WKB.SYSTEM  
Location      ==>                               Project Number ==>
```

Changes to the two user variables for the example

The chart shows the prior values in the template and the new values for the example. You type over the values in the template to create 2 of the user variables. For example, the value in the template was dataset (44) and you replace the value with Data Set Name.

Value in the Template	Value in the Example
Dataset (44)	Data Set Name
LRGDSN	DSNAME
Option Info.	Location
DULLDAY	LOCAL

Value for the third user variable for the example

In the example, you create a third variable. These values are used for the third variable in the example:

- For the variable name, PROJNO
- For the assignend profile variable, &UTUSER03
- For the VPUT value, VPUT(PROJNO)

Example of Modifying the Template

Overview

This unit contains an example of modifying the template. It shows an example of the template before modification and an example of the template after modification. For information about the characteristics of the example, see “Characteristics of the Example of Modifying the Template” on page 474.

Example of the template before modification

The example shows the template before modification. The highlighting indicates the values that will be modified.

```

)ATTR                                /* USER OPTIONS */
  TYPE(INPUT) CAPS(ON)
@ TYPE(OUTPUT) INTENS(HIGH) CAPS(OFF) JUST(LEFT)
.
.
)BODY EXPAND(\\)
@&SFXSCRNM&SFXTITLE\-\
%Cmd%===>_ZCMD
+
+   Data Set (44)  -==>  LRGDSN          +
+   Option Info.  -==>  DULLDAY          +
?
?   To stop the display of this panel modify member KTWPUTUS in the product
?   panel library. Reference the .RESP line in the )INIT section.

THIS IS A SAMPLE PANEL FOR DEFINING UTILITY VARIABLES UNIQUE TO YOUR
INSTALLATION.

YOU DEFINE YOUR OWN VARIABLE NAME AS YOU WOULD FOR ANY OTHER ISPF PANEL.
YOU CAN SET A DEFAULT OR INITIAL VALUE FOR THIS VARIABLE IN THE )INIT
SECTION. THE VALUE MUST BE A 1-32 CHARACTER STRING.
.
.
)INIT
.CURSOR = LRGDSN
.HELP = KTWHTUS
.
.
IF (&UTUSRMSG = &Z)
/* UNCOMMENT THE FOLLOWING LINE TO HAVE THE USER VARS. PANEL DISPLAY
/* &UTUSRACT = YES */
.
.
&UTUSER01 = 'LRGDSN ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */
&UTUSER02 = 'DULLDAY ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */
/* REMOVE COMMENTS TO SET INIT VALUE */
/* IF ( &LRGDSN = ' ',&Z)
/* &LRGDSN = 'TDDB.DBTOOLS.DB2.DEMO.DSWMQ.WKB.SYSTEM' */
)PROC
VPUT ( LRGDSN ) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */
VPUT ( DULLDAY ) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */
)END

```

Example of the template after modification

The illustration shows the template after modification.

```

)ATTR                                /* USER OPTIONS */
  TYPE(INPUT) CAPS(ON)
  @ TYPE(OUTPUT) INTENS(HIGH) CAPS(OFF) JUST(LEFT)
  .
  .
)BODY EXPAND(\\)
c&SFVSCRNM&SFXTITLE\-\
%Cmd%====>_ZCMD
+
+   Data Set Name  ^====>_DSNAME
+   Location       ^====>_LOCAL      +   Project Number ^====>_PROJNO ] 1

)INIT ] 2
.CURSOR = DSNAME ] 3
.HELP = KTWHTUS ]
.
.
IF (&UTUSRMSG = &Z)
/* UNCOMMENT THE FOLLOWING LINE TO HAVE THE USER VARS. PANEL DISPLAY */
&UTUSRACT = YES ] 4
.
.
&UTUSER01 = 'DSNAME ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */
&UTUSER02 = 'LOCAL ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */
&UTUSER03 = 'PROJNO ' /* REQUIRED TO TELL GENERATOR ABOUT VAR */ ] 5
IF (&DSNAME = ' ',&Z)
&DSNAME = 'TDDB.DBTOOLS.DB2.DEMO.DSWMQ.WKB.SYSTEM' ] 6
)PROC
VPUT (DSNAME ) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */
VPUT (LOCAL ) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */
VPUT (PROJNO ) SHARED /* REQUIRED TO PASS VAR TO GENERATOR */ ] 7
)END

```

- 1** Provides instructions and the field for entry for the new user variable
- 2** Marks the end of the text the system displays on the panel (The word *INIT* is on line 39.)
- 3** Specifies the position of the cursor to be in the first field for entry
- 4** Specifies the line as a comment; required to display the panel
- 5** Assigns the third user variable to the profile variable UTUSER03
- 6** Specifies a default value for the Data Set Name field
- 7** Places the third user variable in the shared profile pool

Introduction to Creating Help for User Variables

Overview

This unit contains introductory information about creating Help for user variables. It provides information about the templates to use when creating Help for user variables.

Background about the templates provided within !DB/WORKBENCH

!DB/WORKBENCH provides 3 templates for you to use if you want to create online Help for the user variables you create. The templates include values for these types of items:

- Color
 - The ←PREV and MORE→ reference phrases you use to access the previous Help panel and the next Help panel (if applicable)
 - Reference phrases you use to access standard types of information available in !DB/WORKBENCH Help
-

Resources available with the templates

!DB/WORKBENCH provides 3 templates for you to use to create online Help for the user variables you create. Each template includes the instructions to follow to modify the template you are using. For example, if you want to create Help for one user variable and you need 1 panel, you can follow the instructions in the first template.

Prerequisites for Creating Help for User Variables

Overview

This unit provides information about the prerequisites required for creating Help for user variables.

Knowledge required when creating Help for user variables

To create Help for user variables, you must have experience with ISPF panels. You also must be familiar with the ISPF editor. If you need information, use the IBM documentation on ISPF.

Authority required to create user variables

You must have these authorities to create Help for user variables:

- Authority to update the panel library for !DB/WORKBENCH
 - Authority specific to your site (such as RACF or ACF2)
-

Panel library for the Help panels for user variables

The templates for the Help panels for user variables are in the panel library for !DB/WORKBENCH. The Help panels are in these members:

- KTWHUTUS
 - KTWHUTU2
 - KTWHUTU3
-

Background about the order of the templates

When you create Help for user variables, you must use the template in the member KTWHUTUS first. If you need more than one template for the Help for user variables, follow the instructions in the template in the member KTWHUTUS.

Reminder about protecting the modifications

If you need a reminder about protecting your modifications, see the unit “Prerequisites for Creating User Variables” on page 468.

Background about Templates for Help for User Variables

Overview

This unit provides background information about the templates for Help for user variables.

Elements of the template

The illustration shows the elements in the template that you use when you create Help for user variables.

```

)ATTR DEFAULT((+0)
( TYPE(TEXT) INTENS(HIGH) COLOR(WHITE) /*Standard Text
/* + TYPE(TEXT)INTENS(LOW) COLOR(TURQ) /*Border Infomation
% TYPE(TEXT) INTENS(LOW) COLOR(GREEN) /*Screen Sample
€ TYPE(TEXT) INTENS(HIGH) COLOR(GREEN) /*Labels
@ TYPE(INPUT) INTENS(HIGH) CAPS(ON) COLOR(YELLOW) /*Hot Spot/Command Entry
TYPE(TEXT) INTENS(LOW) CAPS(ON) COLOR(YELLOW) HILITE(REVERSE) /*Warning
)BODY EXPAND(\)
+Cmd ===>@OPT
(\- HELP FOR USER VARIABLES \-
+Press END to return, ENTER+to continue
+or+TAB+to highlighted text and press ENTER to display more information.
€
€Description {This Help panel is a template that you can use to
€ {create the Help for the user variables you create.
€ {It contains Help for the samples on the panel for
€ {user variables.
€
€More Information If you want to know more about creating Help for the
€ {user variables you create, see the panel KTWHTUTS
€ {or the IDB/WORKBENCH User's Guide.
€
€Data Set (44) {The name of the data set you want to use.
€
€Option Info. {The name of your location
€
+ \-
@GLOSSARY @INDEX @COMDS @FLTR @WINDOW @SELECT @MORE ]
)INIT ]
/* Instructions for Creating Help for the panel for User Variables */
/* If you want to use the format used in the Help for IDB/WORKBENCH, */
/* follow these instructions. Type the name of the user variable in */
/* the left column. Type the description for the user variable in the */
/* right column. If you do not want a second Help panel, delete the */
/* @MORE from line 31 and save the changes. If you want a second Help */
/* panel, save the changes and modify the template in the member */
/* KTWHTUT2.
&GLOSSARY = 'Glossary'
&INDEX = 'Index'
&COMDS = 'Commands'
&FLTR = 'Filters'
&WINDOW = 'Windows'
&SELECT = 'Selects'
&PREV = '<-PREV'
&MORE = 'MORE->'
&FILT = 'Filter'
.CURSOR = OPT
)PROC
&CSR = .CURSOR
&ZSEL = TRANS(&CSR
GLOSSARY,*KTWHGLA1
INDEX,*KTWHINDX
COMDS,*KTWHAG20
FLTR,*KTWHAB30
WINDOW,*KTWHMNDW
SELECT,*KTWHAG21
PREV,*KTWHUTUS
MORE,*KTWHUTU2
)
)END

```

Elements of the template (continued)

The list defines the elements of the template.

- 1** Controls the characteristics of the panel
- 2** Provides information and descriptions for the user variables that are the samples on the panels for user variables; serves as a location for the descriptions for your user variables and any instructions you want to provide (You can use lines 14 to line 29 in the ISPF editor for your text.)
- 3** Provides the reference phrases used to access standard types of information available in !DB/WORKBENCH Help; accesses the next Help panel (MORE→) or the previous Help panel (←PREV)
- 4** Marks the end of the text that the system displays on the panel and the beginning of instructions for modifying the template (To execute the panel successfully, be sure that the word *INIT* is on line 30 or above.)
- 5** Provides instructions for modifying the template
- 6** Controls access to standard types of information in the !DB/WORKBENCH Help

Characteristics of the Example of Creating Help for User Variables

Overview

This unit contains information about the characteristics of the example of creating Help for user variables. It includes an illustration that shows the appearance of the Help panel for the example.

Overview of the tasks in the example of creating Help for user variables

These characteristics exist for the example of creating user variables:

- You use the template in the member KTWHUTUS to create Help for the user variables.
- You provide instructions for using the panel for user variables.
- You modify the template to provide descriptions for the three user variables in the example used throughout this chapter.
- You do not need the second or third panel provided with !DB/WORKBENCH.

Appearance of the Help panel for the example

The illustration shows the appearance of the Help panel using the example used in the chapter.

```
Cmd ==>

----- HELP FOR USER VARIABLES -----
Press END to return, ENTER to continue
or TAB to highlighted text and press ENTER to display more information.

Description      Provide the values for the user variables.  The user
                  variables are used in the utility profiles at your
                  site.

Data Set Name    The name of the data set you want to use

Location         The location of your site

Project Number   The number of your project

-----
GLOSSARY  INDEX  COMMANDS  FILTERS  WINDOWS  SELECTS
```

Example of Creating Help for User Variables

Overview

This unit contains an example of an illustration that shows the appearance of the template for Help for the example before modification and the appearance of the template after modification.

Example of the template for Help before modification

The example shows the template for Help for user variables before modification. The highlighting indicates the values that will be modified.

```

)ATTR DEFAULT((+0)
{ TYPE(TEXT) INTENS(HIGH) COLOR(WHITE) /*Standard Text
/* + TYPE(TEXT) INTENS(LOW) COLOR(TURQ) /*Border Infomation
% TYPE(TEXT) INTENS(LOW) COLOR(GREEN) /*Screen Sample
€ TYPE(TEXT) INTENS(HIGH) COLOR(GREEN) /*Labels
@ TYPE(INPUT) INTENS(HIGH) CAPS(ON) COLOR(YELLOW) /*Hot Spot/Command Entry
TYPE(TEXT) INTENS(LOW) CAPS(ON) COLOR(YELLOW) HILITE(REVERSE) /*Warning

)BODY EXPAND(\)
+Cmd ==>@OPT {
{\- \ HELP FOR USER VARIABLES \- \
+Press END to return, ENTER to continue
+or+TAB+to highlighted text and press ENTER to display more information.
€
€ Description {This Help panel is a template that you can use to
€ create the Help for the user variables you create.
€ It contains Help for the samples on the panel for
€ user variables panel.
€
€ More Information{If you want to know more about creating Help for the
€ user variables you create, see the panel KTWHTU5
€ for the IDB/WORKBENCH User's Guide.
€
€ Data Set (44) {The name of the data set you want to use
€
€ Option Info. {The name of your location
€
+ \- \
@GLOSSARY @INDEX @COMDS @FLTR @WINDOW @SELECT @MORE
)INIT
/* Instructions for Creating Help for the panel for User Variables */
/* If you want to use the format used in the Help for IDB/WORKBENCH, */
/* follow these instructions. Type the name of the user variable in */
/* the left column. Type the description for the user variable in the */
/* right column. Save the changes you have made. */
&GLOSSARY = 'Glossary'
&INDEX = 'Index'
&COMDS = 'Commands'
&FLTR = 'Filters'
&WINDOW = 'Windows'
&SELECT = 'Selects'
&PREV = '<-PREV'
&MORE = 'MORE->'
&FILT = 'Filter'
.CURSORS = OPT
)PROC
&CSR = .CURSOR
&ZSEL = TRANS(&CSR
GLOSSARY,*KTWHGLA1
INDEX,*KTWHINDX
COMDS,*KTWH020
FLTR,*KTWH030
WINDOW,*KTWHWINDW
SELECT,*KTWH021
PREV,*KTWHUTU2
MORE,*KTWH0000
)
)END
    
```

Example of the template for Help after modification

The illustration shows the template for Help after modification.

```

)ATTR DEFAULT({+@)
{ TYPE(TEXT) INTENS(HIGH) COLOR(WHITE) /*Standard Text
/* + TYPE(TEXT)INTENS(LOW) COLOR(TURQ) /*Border Infomation
% TYPE(TEXT) INTENS(LOW) COLOR(GREEN) /*Screen Sample
¢ TYPE(TEXT) INTENS(HIGH) COLOR(GREEN) /*Labels
@ TYPE(INPUT) INTENS(HIGH) CAPS(ON) COLOR(YELLOW) /*Hot Spot/Command Entry
TYPE(TEXT) INTENS(LOW) CAPS(ON) COLOR(YELLOW) HILITE(REVERSE) /*Warning

)BODY EXPAND(\)
+Cmd ==>@OPT {
{\- \ HELP FOR USER VARIABLES \- \
+Press END to return, ENTER+to continue
+or+TAB+to highlighted text and press ENTER to display more information.
¢
¢Description {Provide the values for the user variables. The user } 1
¢ {variables are used in the utility profiles at your
¢ {site.
¢
¢Data Set Name {The name of the data set you want to use }
¢
¢Location {The location of your site } 2
¢
¢Project Number {The number of your project
¢
¢
+ \- \
@GLOSSARY @INDEX @COMDS @FLTR @WINDOW @SELECT] 3
)INIT ] 4
&GLOSSARY = 'Glossary'
&INDEX = 'Index'
&COMDS = 'Commands'
&FLTR = 'Filters'
&WINDOW = 'Windows'
&SELECT = 'Selects'
&PREV = '<-PREV'
&MORE = 'MORE->'
&FILT = 'Filter'
.CURSOR = OPT
)PROC
&CSR = .CURSOR
&ZSEL = TRANS(&CSR
GLOSSARY,*KTWHGLA1
INDEX,*KTWHINDX
COMDS,*KTWHAQ20
FLTR,*KTWHAQ30
WINDOW,*KTWHWNDW
SELECT,*KTWHAQ21
PREV,*KTWHUTU2
MORE,*KTWH0000
)
)END

```

- 1** Provides instructions for the user using the panel for user variables
- 2** Provides descriptions of the values the user enters for the user variables
- 3** Specifies that there is not a previous Help panel or information on the next Help panel
- 4** Marks the end of the text the system displays on the panel (The word *INIT* is on line 30.)

Completing the Process of the Example

Overview

This unit contains information on completing the process of creating user variables. It uses the examples used in this chapter.

Modification to the skeleton

To use the user variables in the examples, you also define these variables in the appropriate skeleton:

- &UTUSER01
- &UTUSER02
- &UTUSER03

Running Batch and Using JCL

Chapter 20.

Introduction to !DB/WORKBENCH JCL and Batch

Introduction

This chapter provides an introduction to JCL and batch facilities provided by !DB/WORKBENCH. It introduces actions you must perform if you want to begin to use batch processing and provides background. This chapter is prerequisite to the other chapters in this section.

Chapter contents

Resources for Learning Topics in This Section	490
Overview of the Chapter	492
Introduction to the Process of Using Batch	493
Prerequisites for Using Batch Processing	494
Reports Produced by the Batch Utility	495
Determining What to Do Next	497

Resources for Learning Topics in This Section

Overview

This unit lists resources available in this guide for learning about JCL and batch facilities with !DB/WORKBENCH. It includes information about resources in this guide and in !DB/WORKBENCH.

Background about JCL and batch facilities provided by !DB/WORKBENCH

If you want to use !DB/WORKBENCH online, !DB/WORKBENCH provides facilities so that you can begin to use DB2 utilities and options without additional preparation. You also can choose to customize the skeletons distributed with !DB/WORKBENCH to match the needs of your organization.

If you want to use !DB/WORKBENCH in batch, !DB/WORKBENCH provides an extensive batch facility. You can create batch jobs using commands and selects that you use regularly online and can run them using the batch utility provided by !DB/WORKBENCH.

Benefits of using the batch utility

In !DB/WORKBENCH, you can create and submit batch jobs under ISPF that enable you to perform these types of tasks:

- Duplicate online functions and automate tasks for managing DB2 objects
- Automate operations, perform operations offline (such as running DB2 utilities), and produce reports using functions available with !DB/WORKBENCH

If you want to use !DB/WORKBENCH online

If you want to use !DB/WORKBENCH online, you do not need to use the information in this section. If you want to customize the batch procedure or the skeletons provided by !DB/WORKBENCH, you can follow the instructions in “JCL for Utilities, Batch Jobs, and Extracts” on page 521.

Prerequisites for using this section

This section does not cover details about commands and selects that are available online and supported in batch. It also does not include information about DB2 utilities and options that are supported by !DB/WORKBENCH or about using utilities or utility profiles. Before running jobs in batch, be sure to review the types of information available in the section “Running Utilities” on page 337.

Resources in this section

The chart lists resources within this section, “Running Batch and Using JCL” on page 487.

Information You Want	Resources Available
Introduction to the !DB/WORKBENCH JCL and batch facilities, including prerequisites you must need if you want to use !DB/WORKBENCH in batch	This chapter
Introduction to the requirements you must meet when creating your own batch jobs	“Requirements for Creating Batch Jobs” on page 499
Introduction to and customization of the batch procedure and the skeletons for online or batch processing	“JCL for Utilities, Batch Jobs, and Extracts” on page 521
A dictionary of commands specific to batch processing or the batch utility and commands that operate differently in batch from online	“Command Dictionary for Batch” on page 539

Overview of the Chapter

Organization of information in the chapter

This chapter reviews the prerequisites for batch processing and provides background about how the batch utility handles reports. It covers the basic information you need to begin experimenting with batch processing.

Organization of information and your needs

Review the chart to select information appropriate for your needs.

	Process	Prerequisites	Batch Reports	Determining What to Do
	p.492	p.494	p.495	p.497
Review background about the process of using batch jobs and prerequisites	√	√		
Review background about how !DB/WORKBENCH works with output for reports from such commands as the DUMP, FLIST, and PRNT commands			√	
Determine how to safely experiment or to begin using batch jobs as part of your operations				√

Introduction to the Process of Using Batch

Overview

This unit provides information about the batch procedure and describes the process of using !DB/WORKBENCH batch facilities.

Background about the KTWBATCH procedure

!DB/WORKBENCH provides the KTWBATCH procedure specifically for batch. This procedure contains JCL that handles substitution of variables and executes the batch utility. The KTWBATCH procedure also contains sample batch jobs. To operate !DB/WORKBENCH in batch, build a batch job from the sample commands contained in the KTWBATCH procedure by modifying it for your needs.

Background about the !DB/WORKBENCH skeletons

!DB/WORKBENCH provides many skeletons to manage utilities. You can tailor these skeletons to meet the needs of your site.

Process of using batch

If you want to use commands in batch, follow these broad steps. (These actions are covered in greater detail throughout the chapter.)

Step	Action
1	Confirm that the value for the extract ID in the KTWBATCH procedure is the one that your batch jobs should use.
2	If necessary, provide values for variables in your DD statements.
3	Modify an existing job in the KTWBATCH procedure or create your own job in KTWBATCH.
4	Submit the job for immediate or deferred execution.

Prerequisites for Using Batch Processing

Overview

This unit reviews general prerequisites for success in batch processing and specific requirements such as requirements at installation.

Prior experience and successful batch processing

For successful batch processing, be sure that you have performed these tasks:

- Be proficient in the operation of !DB/WORKBENCH online before beginning to use equivalent operations such as commands, filters, and selects in batch.
 - Review the profiles you have defined for running the DB2 utilities to ensure they meet your needs.
-

Requirements for specifying system information

To use batch processing, you must specify the following values during the installation process. If you want to change these values while executing !DB/WORKBENCH, access the Housekeeping Menu and select option 1 (Profile Data Administration). On the Profile Data Administration panel, select option 1 (DB2-Subsystem Specific Information 1). Verify or update the following values on the panel that is displayed:

- Name of the dataset containing the DSNZPARMS member
 - Name of the member that contains DSNZPARMS
-

Batch processing and user options for Housekeeping

To produce the results you want with batch processing, review the values set for user options during installation. To do this, access the Housekeeping Menu and select option 1 (Profile Data Administration). On the Profile Data Administration panel, select option 3 (DB/Workbench Configuration Information). You can make changes to the configuration of !DB/WORKBENCH on this panel.

Reports Produced by the Batch Utility

Overview

This unit provides background about the reports produced by the batch utility supported by !DB/WORKBENCH. The unit is needed if you plan to use such commands as DUMP, FLIST, or PRNT to produce reports.

Background about the reports provided by the batch utility

When you use the batch utility, it can direct reports to be printed in three ways. If logical separation is needed use the reports DD statement. SYSPRINT is the default destination. Use LISTCPRT to direct AMS LISTCAT output.

DD Statement Used to Generate Report Output	Content of the Report
SYSPRINT	Summary of the commands you used, any error messages produced, and a page index to the REPORTS DD statement report
REPORTS	All the commands you used, as well as output from the DUMP, FLIST, or PRNT commands (if used)
LISTCPRT	Output from the AMS LISTCAT command You must use this DD statement to direct the IDCAMS output to a specified destination (not SYSPRINT). Using this DD statement allows you to receive non-corrupted IDCAMS output.

Background about printing the reports

You can print the reports; however, they can be extensive. For example, a report listing commands used to print all objects, especially printing of authorizations, can result in over a thousand pages, depending on the contents of your DB2 subsystem.

Sample of the SYSPRINT report

The illustration shows a sample SYSPRINT report. In this case, it shows a set of commands that prints all objects. It shows one output line for each input command.

1	PAGE 1	DB	DATA BASES
	PAGE 1	PRNT 2	ALL DATA BASES
		END	
	PAGE 9	TS	TABLE SPACES
	PAGE 9	PRNT	ALL TABLE SPACES
		END	
	PAGE 34	TB	TABLES
	PAGE 34	PRNT	ALL TABLES
		END	
1	PAGE 64	IX	INDEXES
	PAGE 64	PRNT 2	ALL INDEXES
		END	
.			
.			
.			
	PAGE 492	AU D	DATA BASE AUTHORIZATIONS
	PAGE 492	PRNT	PRINT ALL DATA BASE AUTHORIZATION
		END	
	PAGE 522	AU S	TABLE SPACE AUTHORIZATIONS
	PAGE 522	PRNT	
		END	
	PAGE 523	AU T	TABLE AUTHORIZATIONS
	PAGE 523	PRNT	
		END	
	PAGE 903	AU C	TABLE COLUMN AUTHORIZATIONS
	PAGE 903	PRNT	
		END	
1	PAGE 1031	AU V	VIEWS AUTHORIZATIONS
	PAGE 1031	PRNT	
		END	
	PAGE 1056	AU P	PLANS AUTHORIZATIONS
	PAGE 1056	PRNT 2	
		END	
	PAGE 1457	AU B	BUFFERPOOLS AUTHORIZATIONS
	PAGE 1457	PRNT	
		END	
	PAGE 1458	AU G	STORAGE GROUPS AUTHORIZATIONS
	PAGE 1458	PRNT	
		END	
	PAGE 1475	AU U	USERID AUTHORIZATIONS
	PAGE 1475	PRNT	
		END	
	PAGE 1480	END	

1 Example of the page numbers in the REPORT DD statements report for a command that specifies one of the objects

2 Example of the page numbers in the REPORT DD statements report for the PRNT command

Determining What to Do Next

Overview

You can begin to experiment with batch processing or you can begin to use batch processing in your daily operations. This unit reviews different approaches.

If you have never used batch processing with !DB/WORKBENCH

If you have never used batch processing with !DB/WORKBENCH, consider these approaches:

- Experiment with the sample batch commands in KTWBATCH and modify them by adding filters or adding other commands.
 - If you are planning to generate utilities in batch, begin by using the utilities and the utility profiles online; then create your own utility profile and experiment with it online.
-

If you are ready to use batch jobs now

If you are ready to use batch processing, consider these approaches:

- Review the issues in “Requirements for Creating Batch Jobs” on page 499 to confirm that your batch job meets all requirements.
- Review “Determining What to Do Next” on page 520 in that same chapter to confirm that the online commands that you want to use do not require operands or parameters in batch.
- Follow the instructions for customizing the KTWBATCH procedure and the batch skeletons in “JCL for Utilities, Batch Jobs, and Extracts” on page 521 and for protecting your customized versions of these skeletons.

Determining What to Do Next

Chapter 21. Requirements for Creating Batch Jobs

Introduction

This chapter covers the requirements you must meet if you want to create your own batch jobs. It provides essential background and examples ranging in difficulty. The chapter “Introduction to !DB/WORKBENCH JCL and Batch” on page 489 is a prerequisite for using this chapter.

Chapter contents

Overview of the Chapter	500
Methods to Access Object Lists in Batch	502
Methods to Use Commands in Batch	503
Operation of the END Command in Batch	504
Method to Use the PRNT Command to See Reports of Data	506
Methods to Filter Objects in Batch	507
Availability of General Online Functions in Batch	510
Requirements for Using Commands in Batch Jobs	512
Example of a Simple Batch Job	513
Example of a Simple Batch Job for Utilities	514
Example of a Job Stream with Multiple Batch Jobs	515
Background about Commands to Defer Execution	516
Example of a Simple Job for Deferred Execution	518
Example of Multiple Jobs for Deferred Execution	519
Determining What to Do Next	520

Overview of the Chapter

Organization of information in the chapter

The chapter focuses on the requirements for a batch job. It covers background required for creating your own batch jobs and provides examples (with explanations) from a simple job to a deferred execution of multiple jobs.

Information not covered in this chapter

This chapter does not cover information about online functions available with !DB/WORKBENCH. It is intended to build on your existing experience with !DB/WORKBENCH by covering primarily those issues that are different in batch from online. If you are unfamiliar with an online function and want to know more about it, you can find the source in reminders about information covered previously. The units in this chapter also contain references to additional information available in the section “Online Commands and Selects Available in Batch” in *!DB/WORKBENCH Reference Manual*. “Online Commands and Selects Available in Batch” in *!DB/WORKBENCH Reference Manual* contains information for people who prefer information organized as a quick reference.

Relationship between this chapter and other chapters in this section

Before you use the units in this chapter, review the chapter “Introduction to !DB/WORKBENCH JCL and Batch” on page 489. In particular, review the preparation required for using your own batch job.

After you complete the units in this chapter, the commands or selects you want to use in your job determine whether you need to review specific information in “Command Dictionary for Batch” on page 539. Once you create your batch job, follow the instructions in “JCL for Utilities, Batch Jobs, and Extracts” on page 521 to perform necessary preparation of the batch utility and skeletons.

Organization of information and your needs

Review the chart to select information appropriate for the task you want to perform.

	Object List	Use of Commands	END Command	PRNT Command	Filters	General Functions	Requirements	Example—One Job	Example—Utilities	Example—Two Jobs	Control of Execution	Example—Deferred, One Job	Example—Deferred, Two Jobs	Determining What to Do
	p.502	p.503	p.504	p.506	p.507	p.510	p.512	p.513	p.514	p.515	p.516	p.518	p.519	p.520
Review how to use online function in batch and to avoid problems in batch.	√	√	√	√	√	√								
Review requirements for column position, comments, and operands and parameters.							√							
Review rules for using utilities, sorts, and selects in batch.						√								
Review specific information about deferring execution of batch jobs.											√	√	√	
Review specific information about RI data in batch.								√						
Review specific information about utilities in batch.						√			√				√	
Create a simple batch job for the first time.	√	√	√	√	√	√	√	√	√	√				
Determine whether the commands or selects you want to use in your batch job are available in batch or if they have special requirements in batch.														√

Methods to Access Object Lists in Batch

Overview

This unit covers methods for accessing object lists in batch. It includes background about equivalent online operations and resources for determining the correct commands.

Background about equivalent functions available online

When you use !DB/WORKBENCH online to display or to perform operations on objects, you first specify the object by using options available on these menus:

- !DB/WORKBENCH Primary Menu
 - Authorizations Menu (You can access this menu by using the **Z** option on the !DB/WORKBENCH Primary Menu.)
-

Equivalent methods to access object lists in batch

The chart lists the task you want to perform online and the method you can use in batch.

Task You Want to Perform	Action in Batch
Access object lists available if you are using the !DB/WORKBENCH Primary Menu online	Use specific commands available in batch. (For example, use DB for all databases or TS for all table spaces.)
Access authorizations available if you are using the Authorizations Menu online	Use the AU command followed by the letter for the object you want. (For example, use AU D for databases or AU S for table spaces.)

Resources in this section

These resources are available for accessing object lists:

- For a list of the commands available to access object lists in batch, see “Command Dictionary for Batch” on page 539.
- For a list of the operands for the AU command, see “General Commands for Specifying Authorizations in Batch” on page 544 in the chapter “Command Dictionary for Batch.”

Methods to Use Commands in Batch

Overview

This unit compares how you use commands in online and in batch.

Background about equivalent function online

In most cases, when you use a command in !DB/WORKBENCH, you type the command on the command line and press Enter. (When you are using !DB/WORKBENCH online, the system executes the command only if you press Enter.)

Equivalent method to use commands in batch

In batch, listing each command on a separate line in the job is equivalent to typing the command online and pressing Enter. The chart shows lines from a simple batch job and how those lines are equivalent to actions online.

Line From a Batch Job	Your Equivalent Actions Online
TB	On the !DB/WORKBENCH Primary Menu, type T in the Option field and press Enter.
TB=TEST01	On the Tables panel, type TB=TEST01 on the command line and press Enter.
OUT	On the Tables panel, type OUT on the command line and press Enter. On the Process Member menu, type 1 in the Option field and press Enter.
END	On the Tables panel, type END on the command line or press the appropriate function key.

Operation of the END Command in Batch

Overview

This unit shows examples of the three purposes of the END command in batch. It includes background about the END command online.

Background about specific functions of the END command in batch

In batch, you use the END command as the last command in the batch job. The END command serves to terminate the batch utility.

Background about equivalent functions available in batch

When used online, the END command returns to any prior panel. For example, you could use the END command in these ways online:

- To return to the original object list panel after using a command that displays data about an object
- To return to the !DB/WORKBENCH Primary Menu because you have completed all tasks you want to perform with the type of object

Background about differences in the END command online and in batch

When you use more than one of the commands that are most often followed by the END command, you use the END command differently online than you do in batch. This chart shows the differences.

Type of Processing You Are Using	Requirement for the END Command	Example of Using Multiple Commands
Using a command that requires the END command in online processing	Use the END command after each of these commands. (You can, however, do other actions, such as printing, before using the END command.)	Your actions online: Use DEP Use PRNT Use END Use DBR Use PRNT Use END
Using a command that requires the END command in batch processing	Use the END command one time, no matter how many of this type of command you have in your job.	Your batch commands: DEP PRNT DBR PRNT END

Equivalent methods in batch

In batch, you use the END command for the same purpose for which you use the command line online. The examples reflect these common purposes.

Example 1

```
TB
TSSET
PRNT
END 1
END 2
END 3
```

Example 2

```
SG
VOLS
PRNT
END 1
END 2
END 3
```

Example 3

```
PL
COLL
PRNT
END 1
END 2
END 3
```

- 1** Completes a command that requires the use of the END command when used online (Equivalent to the action online of using the END command on the panels displayed when you use the TSSET, VOLS, or COLL commands.)
- 2** Indicates that no other command in the batch job applies to this type of object (Equivalent to the action online of using the END command on the Tables, Storage Group, or Plans panels to return to the !DB/WORKBENCH Primary Menu.)
- 3** Terminates the batch utility

Method to Use the PRNT Command to See Reports of Data

Overview

This unit covers how to use the PRNT command to see reports of data.

Background about equivalent functions available online

When you display object lists online with !DB/WORKBENCH, you can review the objects displayed or you can choose to use the PRNT command if needed. The PRNT command, when used online, sends output to the system printer.

Equivalent methods to see reports of data in batch

In batch, if you want to see the object lists that result from a command, you must also use the PRNT command for each report of data you want to see. The PRNT command, when used in batch, sends output either to the SYSPRINT report or the REPORT DD statement report. (The location of the output varies with your input job stream.)

Examples of using the PRNT command

The chart shows examples of how to use the PRNT command to generate batch reports for the object lists you want to be able to review.

Object List You Want to Print	Lines from the Batch Job
All storage groups	SG PRNT END
All storage groups by volume	SG VOLS PRNT END
All storage groups and all storage groups by volume	SG PRNT VOLS PRNT END

Methods to Filter Objects in Batch

Overview

This unit provides suggested actions you can perform using filtering to prevent errors and provides examples and resources for filtering.

Background about the importance of filtering in batch

The filters you use can make the difference between success and failure of batch processing. The chart shows two examples of problems and suggested preventive actions.

Example of Problem	Suggested Preventive Actions
<p>In your batch job, you use a DB filter and then later a CR filter. When used together, the two filters results in no objects meeting the criteria. You subsequently use the OUT command in the batch job.</p> <p>When you check, you discover that no CREATE statements were generated.</p>	<p>Use the NOFILT (or NOF) command in the batch job before a filter to ensure that filters you set previously do not affect the current job stream.</p> <p>Experiment online to be sure that any filters you want to use together result in the objects you want.</p>
<p>In your batch job, you use a filter that results in such a large list of objects that the time allotted for your job is less than the required time to run the batch job to completion.</p>	<p>Apply additional filters to shorten the object list to be processed in batch.</p> <p>Increase the time allotted to your job.</p>

Background about filtering and batch

You can use in batch most of the methods for filtering that you can use online from the command line. The exception is that you cannot use two filters combined with a semicolon (;).

Examples of filtering of names in batch

The chart shows examples of filtering names in batch. (You also can use the asterisk when providing the database name with the FLIST command.)

Objects You Want	Lines from a Batch Job
Tables associated with the database <i>TST1</i> and any of the table spaces that begin with the characters <i>TESTW</i>	TB NOFILT DB=TST1 TS=TESTW*
Databases with the same three initial characters (<i>TST</i>) and last characters (<i>961</i>)	DB NOFILT DB=TST**961

Examples of EQUAL vs. NOT EQUAL filtering command in batch

The chart shows examples of using the NOT EQUAL filter in batch compared to using an EQUAL filtering command in batch.

Objects You Want	Lines from a Batch Job
All databases that do not begin with <i>DSN</i> (i.e., databases not equal <i>DSN</i> in first three characters of name) (PRNT command prints selected databases to location specified in input job stream.)	DB DB=DSN* FILT=N PRNT END
Only databases that begin with <i>DSN</i> (i.e., databases equal <i>DSN</i> in first three characters of name) (PRNT command prints selected databases to location specified in input job stream.)	DB DB=DSN* PRNT END

Examples of filtering of values or statuses in batch

The chart shows examples of filtering values or statuses in batch.

Objects You Want	Lines from a Batch Job
All tables with an audit trail for changes only	TB NOFILT AUDIT=C
All plans using the EXPLAIN option	PL NOFILT EXPLAIN=Y
All authorizations for plans with authority to copy a package	AU P NOFILT CO=G

Examples of filtering numeric values

The chart shows examples of filtering numeric values, such as the quantity of an object (for example, the number of views) or a number distinctive for an object (for example, the maximum length for a record). When you filter numeric values, you can use =, ≠, <, >, <=, or >= as operands.

Objects You Want	Lines from a Batch Job
All plans with no DBRMs	PL NOFILT DBRMS=0
All databases with less than 5 table spaces	DB NOFILT SPACES<5
All tables that have relationships defined for referential integrity	TB NOFILT REFS>0

Reminder about filtering

If you need a reminder about how to use filters, see “Filtering Object Lists” on page 195.

Additional information about filtering

For additional information about filtering, see the chapter “Filtering Object Lists” in *!DB/WORKBENCH Basic User's Guide*.

Availability of General Online Functions in Batch

Overview

This unit reviews the availability of frequently used online functions and resources if you need additional information. It also covers a variation in the operation of the S command in batch from its operation online.

Background about functions used from the command line and batch

In most cases, you can use in batch most functions that you can use online from the command line. You cannot, however, use the GEN command, which is used within the utilities, or commands that display selection lists such as the dash (-), FILT, DOP, DOS, and DO commands. You also cannot use the slash (/) command used to create a window to access other object lists.

Functions for utilities in batch

You can perform in batch most functions for utilities that you can perform from the command line. (If you need a reminder about accessing utilities, setting utility profiles, or using the MOD parameter, see “Running Utilities” on page 337.) The chart provides examples of functions for utilities you can use in batch.

Tasks you Want to Perform	Lines from a Batch Job
Set (make active) a utility profile named <i>TEST</i> .	TS UTIL OPTS TEST
Use the REORG utility with the table space named <i>TEST01</i> .	TS NOFILT TS=TEST01 UTIL REORG
Use the REORG and COPY utilities with the table space named <i>TEST01</i> . (You want to combine the generated statements into one member.)	TS NOFILT TS=TEST01 UTIL REORG UTIL COPY MOD

Functions for sorting in batch

The chart provides examples of sorting that you can do in batch. (If you need a reminder about sorting, see the chapter “Using ISPF Facilities within !DB/WORKBENCH” in *!DB/WORKBENCH Basic User's Guide*.)

Tasks you Want to Perform	Lines from a Batch Job
Sort in ascending order all database names.	DB SORT DB
Sort in descending order all storage groups.	TS SORTD SG

Function to use a select in batch

If you want to use a select in batch, you use the S command followed by the letter for the select you want.

Tasks you Want to Perform	Lines from a Batch Job
Use the RI Info (F) select with referential data.	TB NOFILT REFS>0 S F
Use the indexes for the table spaces whose names begin with the letters TEST .	TS TS=TEST* S I

Consequences of the S command in batch

The number of objects that the S command affects in batch is different from the number it affects when it is used online.

- Used online, S affects only the objects currently displayed on the panel
- Used in batch, S affects all objects on the object list

Requirements for Using Commands in Batch Jobs

Overview

This unit contains requirements for using commands in a batch job.

Column position, commands, and comments

These rules apply to column positions:

- Begin a command in any column position; however, the command must be within columns 1-71. (The batch utility ignores data in columns 72 or higher.)
 - Place an asterisk (*) in column 1 to indicate a comment. (The batch utility ignores any data that has an asterisk in column 1.)
-

Operands and parameters used with commands

These rules apply to operands and parameters used with commands.

Format Used for the Command	Rule to Follow	Example
The operand or parameter includes a parenthesis	No spaces between the words and the parenthesis	LIMIT(100) DB(TESTDB) MAC(TESTEDIT)
The operand or parameter includes an equals (=) sign	No spaces between the words and the equals sign	DB=TESTDB DBRMS>1 PDSCOPY, TO=PROD.JCL(JOB02)
The operand or parameter consists of separate words	Spaces between the words	UTIL OPTS TEST UTIL COPY MOD AMS DEFINE ALL

Example of a Simple Batch Job

Overview

This unit shows the elements of a simple batch job.

Elements of the batch job

The example shows essential elements of a sample batch job.

```

1
* RI CHECK INDEX/CHECK DATA
TB 2
NOFILT 3
REFS>0 4
S F }
SET 5
RFO }
END 6
END 7

```

- 1** Optional comment (The asterisk for the comment is in column 1.)
- 2** Required command to specify the object for the commands to act on
- 3** Optional use of the NOFILT command (Using the NOFILT command can eliminate problems with previously set filters.)
- 4** Optional filters (Using filters can avoid errors in processing and speed up time required for processing.)
- 5** Commands for the batch utility to use with the object
- 6** Required use of the END command to indicate the last command for the batch utility to apply to the object
- 7** Required use of the END command to terminate the batch utility

Example of a Simple Batch Job for Utilities

Overview

This unit shows the elements of a simple batch job for utilities.

Elements of the batch job

The example shows essential elements of a sample batch job for utilities.

```
TS 1  
NOFILT  
DB=TEST06  
UTIL OPTS TEST 2  
UTIL COPY 3  
UTIL REORG MOD 4  
END 5  
END 6
```

- 1** Required command to specify the object for the commands to act on
- 2** Optional use of the command to set (make active) the utility profile named *TEST*
- 3** Required command to use the COPY utility
- 4** Optional use of the MOD operand to add the REORG statements to the end of the member
- 5** Required use of the END command to indicate the last command for the batch utility to apply to the object
- 6** Required use of the END command to terminate the batch utility

Example of a Job Stream with Multiple Batch Jobs

Overview

This unit shows the elements of two simple batch jobs used together.

Elements of the batch jobs

The example shows essential elements when you want to use multiple objects.

```
PK
PRNT
DEP
PRNT 1
END
END
PL
PRNT
DBR
PRNT 2
DEP
PRNT
END
END
END 3
```

- 1** First job (beginning with the command for the first object and completing with the END command that completes the DEP command and then with the END command)
- 2** Second job
- 3** Required END to terminate the batch utility

Background about Commands to Defer Execution

Overview

This unit reviews reasons for deferring execution of the batch job and lists commands available to control the execution of the job or to manage the PDS member.

Background about reasons for deferring execution

Among the reasons for deferring execution of commands in batch are these examples.

- Use a package for job scheduling
 - Combine statements from multiple utilities in one PDS member
 - Copy commands to another dataset
-

Commands available to control execution of commands

These commands are available to control the execution of commands or of the contents of a newly created PDS member:

- EXEC command, the default in batch (immediately submits a job or executes the statements; turns off the NOEXEC command)
 - NOEXEC command (defers submission of a job or execution or statements; turns off the EXEC command)
 - PDSEXEC command (executes a newly created member if the NOEXEC command has previously been used; turns off the NOEXEC command)
-

Commands available to manage the PDS member

In addition to the PDSLIST and PDSEXEC commands, the batch utility provides these commands to manage the PDS member created during batch:

- PDSCOPY command (copies the contents of one PDS member to another member, either in the !DB/WORKBENCH user PDS or to a dataset you specify)
- PDSEEDIT command (invokes the ISPF editor and a specified ISPF edit macro)

Resources in this section

If you want to see examples of some of the commands used to defer execution, continue with the units in this chapter. If you want more information about the commands themselves, see the alphabetic entry for the command in “Command Dictionary for Batch” on page 539.

Example of a Simple Job for Deferred Execution

Overview

This unit shows the elements of a job that results in deferred execution. (The unit “Background about Commands to Defer Execution” on page 516 is a prerequisite for this unit.)

Elements of the batch job

The example shows essential elements of the batch job that defer execution.

```
TS 1  
NOFILT  
NOEXEC 2  
DB=TST1  
TS=TESTW*  
UTIL OPTS CANDLE  
UTIL COPY  
PDSCOPY TO=CANDLE.SCHED.TEST(JOB01) 3  
END 4  
END 5
```

- 1** Required command to specify the object for the commands to act on
- 2** Required command to turn off the EXEC command (the default for the batch utility)
- 3** Command to copy the PDS member to another dataset
- 4** Required use of the END command to indicate the last command for the batch utility to apply to the object
- 5** Required use of the END command to terminate the batch utility

Example of Multiple Jobs for Deferred Execution

Overview

This unit shows the elements of multiple jobs that result in deferred execution.

Elements of the batch jobs

The example shows essential elements of the batch jobs that defer execution.

```

TS
NOFILT
NOEXEC 1
DB=CTEST1
TS=TEST2*
2 UTIL OPTS CANDLE
UTIL COPY
UTIL MODIFY MOD
PDSCOPY TO=WKB100.BATJCL2.CNTL(JOB01)
END
TS
NOFILT
NOEXEC 3
UTIL OPTS T00LS
4 UTIL COPY
UTIL MODIFY MOD
UTIL RUNSTATS MOD
PDSCOPY TO=WKB100.BATJCL2.CNTL(JOB02) MOD
END
END 5

```

- 1** Required command to turn off the EXEC command
- 2** First job (beginning with the command for the object and completing with the END command that indicates the last command to act on the object)
- 3** Required command to turn off the EXEC command
- 4** Second job
- 5** Required use of the END command to terminate the batch utility

Determining What to Do Next

Overview

This unit covers the options and resources available to you when you are ready to create your own batch job with commands appropriate for your organization.

Options and resources available

Review the chart for options and resources available.

Task You Want to Perform	Source to Use
Determine whether the commands you want to use are available in batch.	“Online Commands and Selects Available in Batch” on page 89
Use one of these online commands that requires operands or parameters in batch: <ul style="list-style-type: none"> ● AMS command ● BIND command ● DISP command ● OUTC command ● REBIND command ● START command ● UTIL command 	“Command Dictionary for Batch” on page 539 (The chapter is organized alphabetically.)
Use one of these commands, available in batch, that requires additional information: <ul style="list-style-type: none"> ● DUMP command ● FLIST command 	“Command Dictionary for Batch” on page 539 (The chapter is organized alphabetically.)

Chapter 22.

JCL for Utilities, Batch Jobs, and Extracts

Introduction

This chapter contains background about the skeleton JCL and symbolic variables that !DB/WORKBENCH provides for utilities, batch jobs, and extracts. It also describes the KTWBATCH procedure, which is used to perform !DB/WORKBENCH operations in batch.

Chapter contents

Overview of the Chapter	522
Background about Skeletons for Online and Batch	524
How to Customize a Skeleton or the KTWBATCH Procedure	526
Customization of the KTWUTJCL Skeleton	528
Introduction to !DB/WORKBENCH and the Batch Utility	530
Background about Customization of the KTWBATCH Procedure	531
Making Recommended Modifications for PDS Names	532
Making Recommended Modifications for DD Statements	533
Making Recommended Changes to the Extract ID	534
Making Optional Modifications to Batch Jobs	535
Modifying DD Statements for Batch Extracts	537
Submitting a Batch Job for Execution	538

Overview of the Chapter

Overview

This chapter begins with an overview of the !DB/WORKBENCH skeletons and the KTWBATCH procedure and describes the steps to prepare the batch utility and create a batch job. The section describes actions you can take if you want to modify the distributed skeletons, and guidelines for preparing your own batch jobs.

Organization of information and your needs for the batch utilities

If you will use JCL created during online processing to run the batch utilities, review this chart to select information appropriate to your needs.

	Background About Skeletons	Customizing a Skeleton	Customizing the Skeleton Driver	Submission for Execution
	p.524	p.526	p.528	p.538
Review the contents and purpose of the skeletons	√			
Review customization of the driver skeleton for utility generation		√		
Review customization of sizing parameters		√		
Review defining job card information			√	
Review how to submit a job				√

Organization of information and your needs for running batch jobs

If you will use JCL created during online processing to create and run batch jobs, review this chart to select information appropriate to your needs.

	Background About Skeletons	Customizing a Skeleton	Introduction to the Batch Utility	Customizing the Batch Utility	Recommendations for PDS Names	Recommendations for DD Statements	Recommendation for Extract ID	Optional Modifications	Submission for Execution
	p.524	p.526	p.530	p.531	p.532	p.533	p.534	p.535	p.538
Review the contents of the skeletons	√								
Review how to customize a skeleton		√							
Review the contents and use of the KTWBATCH procedure			√						
Review how to customize the batch utility				√					
Review recommended modifications					√	√	√		
Review optional modifications								√	
Review how to submit a job									√

Background about Skeletons for Online and Batch

Overview

This unit provides background about skeletons provided by !DB/WORKBENCH. Skeletons are used by both batch and online processing.

Note: The skeletons provided by !DB/WORKBENCH are samples that may change with maintenance.

Background about the index to skeletons for utility generation

The KTWUT\$IX skeleton, located in the SKELS library, provides an index to the skeletons for utility generation. The chart lists the types of information in the KTWUT\$IX skeleton and provides an example.

Information in the Skeleton	Examples
Name of each of the utility skeletons provided by !DB/WORKBENCH and a brief explanation of its purpose.	KTWUTCKD—CHECKDATA KTWUTSTA—RUNSTATS CALLED BY KTWUTJCL (MAIN DRIVER), KTWUTRGT, (REORG TABLESPACE) AND KTWUTRGX (REORG INDEX)
Name of each ISPF variable (If a variable is new, it is indicated by an asterisk.)	CKCOPY—CHECK DATA COPY EXCEPTION DRIVER *KEEPDICT—Y,N TO REUSE COMPRESSION DICTIONARY
Name of each ISPF space variable and information about primary and secondary sizing and how the variables interact with the parameters for space control in the KTWUTSIZ skeleton	CPSPCP—PRIMARY COPY SIZE SORECS—SECONDARY SORT SIZE—RECS
Change history for the skeleton and variables (if any)	

Background about the driver for skeletons for utility generation

The KTWUTJCL skeleton, located in the SKELS library, is necessary for utility generation in both batch and online. It serves these purposes:

- Drives all of the skeletons for utility generation (such as the skeleton for the CHECKDATA utility)
- Contains control card information and variables for sizing and space control
- Contains statements for running the DB2 utilities in batch
- Contains values read from the current utility profile or from your entry of values online when you select the utility
- Provides explanations of sizing values and space control
- Provides instructions for customizing the skeleton

How to Customize a Skeleton or the KTWBATCH Procedure

Overview

This unit provides background and basic procedures for customization and is applicable either to online or batch.

Background about customization

You can use !DB/WORKBENCH online or in batch without customizing the skeletons in the SKELS library, but you must customize the KTWBATCH procedure in the CNTL library.

Cautions to consider before customization

When you install a new version of !DB/WORKBENCH, the skeletons are replaced in the SKELS library, and KTWBATCH is replaced in the CNTL library. If you have customized these libraries in previous versions, a new installation of !DB/WORKBENCH will overwrite them. To avoid loss of your modifications, copy the members you want to customize and place the modified members in your own libraries.

Requirements for customization

To be successful in customization, meet these requirements:

- Make all modifications in a copy of the skeleton or procedure
- Place each modified skeleton or procedure into a separate member.
- Name the member so that it *does not* begin with the letters *KT*.
- Place your customized skeletons in your own skeleton library.

Background about !DB/WORKBENCH and customization

If you are using !DB/WORKBENCH and modify a skeleton or the KTWBATCH procedure in either your own skeleton or CNTL library or in the distributed skeleton or CNTL library, you must exit and restart !DB/WORKBENCH so that the customization takes effect.

Procedure to follow for successful customization

Be sure you have reviewed requirements for customization. If you choose to customize, follow these steps.

Step	Action
1	Open a member in your own CNTL or SKELS library. (You can use ISPF option 3.1 also.)
2	Identify the KTWBATCH procedure (in the product installed CNTL library) or the skeleton (in the product installed SKELS library) you want to customize. Copy it to the member you opened. (You copy from the CNTL library to your CNTL library and from the SKELS library to your SKELS library.) In this case, the product installed member and the customized member exist in different libraries so they can have the same member name. or You can open a new member in the product installed library (CNTL or SKELS) and customize the new member. In this case, the product installed member and the customized member exist in the same library so they must have different member names.
3	Customize the member as required by your installation.
4	Place the customized member in your own CNTL or skeleton library.
5	Concatenate your CNTL or skeleton library in front of the respective !DB/WORKBENCH CNTL or skeleton library.
6	Change the KTWUTJCL skeleton to point to the member containing the customized skeleton, if you changed the name.
7	Change the startup CLIST KTCSETUP to point to the concatenated libraries.

Customization of the KTWUTJCL Skeleton

Overview

This unit contains essential background if you want to customize the KTWUTJCL skeleton for utility generation. The KTWUT\$IX skeleton, located in the SKELS library, provides an index to the skeletons for utility generation. (The units “Background about Skeletons for Online and Batch” on page 524 and “How to Customize a Skeleton or the KTWBATCH Procedure” on page 526 are prerequisites for this unit.)

Requirement for the skeleton library if you customize

The skeleton KTWUTJCL resides in the SKELS library. To execute the KTWUTJCL skeleton, you must be sure that all of the skeletons that it drives are present either in the SKELS library or in your own skeleton library.

Types of customization you can perform

The KTWUTJCL skeleton makes it possible to readily customize these items:

- Sizing parameters
 - Job statements
-

Background about options for customizing sizing parameters

The KTWUTJCL skeleton contains variables such as SPCCTL, BLKSIZ, and REGION that you can change for your operations. The skeleton contains these items:

- Variables themselves
- Information about possible values and the default values
- Explanation of values and implications of changing some values

Background about options for customizing job statements

You have either of these options for job statements with !DB/WORKBENCH:

- JOB CARD in your own member
- Up to four lines of job statements

Background about methods for defining job statements

Whether you use a JOB CARD in your own member or provide job statements, you define job control on the General Options panel (KTWPUTSV). You can provide the values online or by using a utility profile. Once you have defined the member name for the JOB CARD or the job statements, !DB/WORKBENCH retains the values from session to session. If you use utilities in batch, the KTWUTJOB skeleton provides the values at execution.

Limitation for customizing job statements

Although you can use either a JOB CARD in your own member or up to four lines of job statements, you cannot do both. If you provide both the member name and the job statements, !DB/WORKBENCH ignores the job statements.

Introduction to !DB/WORKBENCH and the Batch Utility

Overview

This unit provides a brief introduction to the broad steps !DB/WORKBENCH performs in batch. It introduces a series of units showing how you can perform recommended modification to the KTWBATCH procedure.

Background about the procedure for batch processing

The KTWBATCH procedure, located in the CNTL library, is required for batch processing. It serves these purposes:

- Substitutes required parameters
 - Executes the batch utility
 - Provides a variety of batch jobs you can either use without additional preparation or modify to meet your needs.
-

Actions performed by !DB/WORKBENCH and the batch utility

!DB/WORKBENCH and the batch utility execute jobs in batch by performing these broad steps.

Step	Action
1	!DB/WORKBENCH reads the JCL in the KTWBATCH procedure to determine these items: <ul style="list-style-type: none"> ● The extract ID to use with this batch job ● The DDnames and profile names taken from the datasets KTCPROFI and KTCPROFO
2	!DB/WORKBENCH reads the profile associated with the extract ID to determine the required dataset names.
3	!DB/WORKBENCH imbeds the skeletons and substitutes the values from the utility profile that is active.
4	!DB/WORKBENCH executes the commands in the batch job you included. (The batch jobs follow the SYSIN DD statement.)
5	!DB/WORKBENCH writes the results to the !DB/WORKBENCH user PDS.

Background about Customization of the KTWBATCH Procedure

Overview

This unit contains essential background for customizing the KTWBATCH procedure and introduces recommended items to customize. (The unit “How to Customize a Skeleton or the KTWBATCH Procedure” on page 526 is a prerequisite for this unit.)

Cautions about customizing the KTWBATCH procedure

These cautions apply to customizing the KTWBATCH procedure:

- If you plan to change KTWBATCH, you should copy and rename it in order to save your changes.
 - If Candle® Corporation distributes changes to KTWBATCH, it is your responsibility to ensure that the changes are made to the copies of KTWBATCH you are maintaining.
-

Determining what to do next

Review the chart to determine the source to use for what you want to do next.

Task You Want to Perform	Source to Use
Review and modify parameters that control batch processing	“Making Recommended Modifications for PDS Names” on page 532 “Making Recommended Modifications for DD Statements” on page 533 “Making Recommended Changes to the Extract ID” on page 534
Customize a batch job or insert or delete a job	“Making Optional Modifications to Batch Jobs” on page 535

Making Recommended Modifications for PDS Names

Overview

This unit is the first in a series of recommended modifications to the KTWBATCH procedure. (The unit “How to Customize a Skeleton or the KTWBATCH Procedure” on page 526 is a prerequisite for this unit.)

Background about modifications for PDS names

To modify the KTWBATCH procedure to provide for PDS names, you replace these statements with the PDS names for your site:

```
//KTWBATCH PROC SOUT '*'  
' CANDLE.USER.LOAD '  
' CANDLE.USER.MSGS '  
' CANDLE.USER.SKELS '  
' CANDLE.USER.PANELS '  
' CANDLE.USER.PROFILE'
```

Making the changes in the KTWBATCH procedure

You can provide the PDS names for the PROC parameters for messages, skeletons, panels, and profile. Follow these steps:

Step	Action
1	Locate the statements in the KTWBATCH procedure
2	Replace the default PDS names with the PDS names appropriate for your site.
3	Perform the appropriate action: <ul style="list-style-type: none">● Save the member.● Make additional modifications at this time.

Making Recommended Modifications for DD Statements

Overview

This unit is the second in a series of recommended modifications to the KTWBATCH procedure. (The unit “How to Customize a Skeleton or the KTWBATCH Procedure” on page 526 is a prerequisite for this unit.)

Background about modifications to DD statements

To modify the KTWBATCH procedure, provide SYS1ISPF PDS names for the DD statements. You replace these statements with the names appropriate for your site.

```
'SYS1.ISPF.LOAD'
'SYS1.ISPF.PANELS'
'SYS1.ISPF.MSGS'
'SYS1.ISPF.TABLES'
```

Making the changes in the KTWBATCH procedure

You can provide the PDS names for the load library, panels, messages, and tables DD statements. Follow these steps:

Step	Action
1	Locate the statements in the KTWBATCH procedure
2	Replace the default PDS names with the PDS names appropriate for your site.
3	Perform the appropriate action: <ul style="list-style-type: none"> ● Save the member. ● Make additional modifications at this time.

Making Recommended Changes to the Extract ID

Overview

This unit is the third in a series of recommended modifications to the KTWBATCH procedure. (The unit “How to Customize a Skeleton or the KTWBATCH Procedure” on page 526 is a prerequisite for this unit.)

Significance of the extract ID in batch

If the extract ID specified in the KTWBATCH procedure is not the one you need, your batch job can produce incorrect results or no results. For example, if the objects you specify to your batch job are not in the extract, !DB/WORKBENCH cannot produce output.

Background about the default extract ID in KTWBATCH

The KTWBATCH procedure includes the ISPSTART statement which specifies the extract ID for the batch utility to use. The default value (which follows the PARM operand) for the extract ID is DB2.

Making the changes in the KTWBATCH procedure

You can provide the extract ID for !DB/WORKBENCH to use when running the batch utility. Follow these steps:

Step	Action
1	Locate the ISPSTART statement in the KTWBATCH procedure
2	Replace the default extract ID (DB2) in the PARM field with the value appropriate for your site.
3	Perform the appropriate action: <ul style="list-style-type: none">● Save the member.● Make additional modifications at this time.

Making Optional Modifications to Batch Jobs

Overview

This unit provides background about modifying a batch job, with examples. It also provides information about the relationship between various !DB/WORKBENCH procedures and skeletons used to run batch.

Background about the groups of jobs

The KTWBATCH procedure contains many examples of commands for such tasks as printing reports about the data. The following example is taken from the KTWBATCH procedure.

```

/*
//SYSIN DD *
PL          PLANS
PRNT       PRINT ALL PLANS
  DEP      PLANS DEPENDENTS
  PRNT     PRINT ALL PLAN DEPENDENTS
  DBR      PLANS DBRMS
  PRNT     PRINT ALL PLANS DBRMS
END        END SUBDISPLAY
DBRMS=0    DISPLAY PLANS WITH NO DBRMS
PRNT
END        END MAJOR DISPLAY

```

Relationship between the KTWBATCH procedure and the KTWUTJCL skeleton

You use the sample commands in the KTWBATCH procedure to print reports for the DB2 objects. KTWBATCH calls the KTWUTJCL skeleton that contains the control statements for running the DB2 utilities.

Reminder about batch jobs

If you need a reminder about the requirements for creating batch jobs, review “Requirements for Creating Batch Jobs” on page 499.

Examples of a modification

The chart below shows examples of batch jobs as they appear in KTWBATCH, and as they appear after they have been modified.

<p align="center">Job as It Appears in KTWBATCH</p>	<p align="center">Modified Job</p>
<p>This job prints all tables.</p> <pre> TB TABLES PRNT PRINT ALL TABLES END </pre>	<p>This modified job filters for all tables starting with the letters <i>TST</i> before printing. A command has been added (1) to perform the filtering.</p> <pre> TB TABLES TB=TST* 1 PRNT PRINT ALL TABLES END </pre>
<p>This job prints all storage groups first in total and then by VOLSER.</p> <pre> SG PRNT 2 VOL PRNT END </pre>	<p>This modified job prints the storage groups by VOLSER only. The command to print all storage groups (2) is removed.</p> <pre> SG VOL PRNT END </pre>

Modifying DD Statements for Batch Extracts

Overview

This unit provides background information and the recommended JCL modification for running large batch extracts that exceed sort capacity.

Background about large batch extracts

Batch extracts are handled differently than online extracts, and large batch extracts may not complete and return a message that the sort capacity has been exceeded.

Making the changes in the DD statements

You can add more sortwork datasets by adding sortwork DD statements to your batch extract JCL. The following statement would create a sortwork dataset using a disk device:

```
//SORTWK01 DD SPACE=(CYL,(15,5)),UNIT=3390
```

Submitting a Batch Job for Execution

Overview

This unit provides information about the types of jobs that can be submitted for batch processing and tells you how to submit the job for processing.

Background about jobs that can be processed in batch

Jobs that can be submitted for batch execution can be one of the following:

- The KTWBATCH procedure
 - JCL created by KTWBATCH processing and stored in a member of your user PDS by a PDSCOPY command (You specify the name of the member in which to store the JCL using a parameter of the PDSCOPY command.)
 - JCL created as output of the online utility generation process (!DB/WORKBENCH stores this JCL in your user PDS in a member named *useridU*.)
-

Steps in submitting a job for execution

Follow the steps in this chart to submit your batch job for execution.

Step	Action
1	Open the member that contains the JCL for your batch job.
2	Type SUBMIT or SUB on the command line.
3	Press Enter. Result: !DB/WORKBENCH submits your batch job to be processed.

Introduction

This chapter covers commands that are available when you use the batch utility (for example, the DUMP command), or that you need to manage the batch utility (for example, the EXEC command). It also contains online commands that operate differently in batch (for example, the BIND command) and commands that are equivalent to online commands (for example, the DB and AU D commands). (The chapter “Introduction to !DB/WORKBENCH JCL and Batch” on page 489 is a prerequisite for this chapter.)

Chapter contents

Overview of the Chapter	540
General Commands for Specifying Objects in Batch	543
General Commands for Specifying Authorizations in Batch	544
AMS Command and Parameters for Batch	545
BIND Command and Parameters for Batch	547
DISP Command and Parameters for Batch	548
DUMP Command and the Batch Utility	549
END Command and Options in Batch	552
EXEC Command	554
FLIST Command	555
NOEXEC Command and the Batch Utility	558
OUTC Command and Parameters in Batch	559
PDSCOPY Command and the Batch Utilities	562
PDSEEDIT Command and the Batch Utility	564
PDSEXEC Command and the Batch Utility	566
PDSLIS Command and the Batch Utility	567
PRNT Command and Operation in Batch	569
START Command and Parameters in Batch	571
UTIL Command and Functions Supported in Batch	572

Overview of the Chapter

Background about online commands and batch processing

In general, the commands you use online work in the same way when you use them in batch. Some of these online commands, however, require parameters when used in batch. Other commands, such as the END and PRNT commands, serve additional functions in batch jobs.

Background about online menus and batch processing

The batch utility provides commands equivalent to the options for accessing objects from the !DB/WORKBENCH Primary Menu and from the Authorizations Menu.

Background about commands specific to batch processing

The batch utility provides commands to control execution of batch jobs, to manage the PDS member, and to perform additional functions.

Background about the UTILCMDS DD statement

The UTILCMDS DD statement lists valid batch commands and associates the batch commands with specific processor modules. The UTILCMDS DD statement is required.

Organization of information in this chapter

With the exception of the first two units that cover general commands to access object lists or authorizations, the units are organized in alphabetical order. The units contain these items:

- Description of the purpose of the command
- Prerequisites (if any) or background about requirements (if any)
- Format for using the command (if applicable)
- Requirements for using the command, including the location of the output from the command (if applicable) and the operands and parameters
- Examples (if appropriate)

Relationship between this chapter and other chapters

This chapter covers only information specific to batch processing. These other resources are available:

- If you want to know whether a command or select is available in batch, see “Online Commands and Selects Available in Batch” on page 89.
- If you want to see how to use these commands in batch jobs, see “Command Dictionary for Batch” on page 539.

Background about symbols used in this chapter

These principles are used in showing syntax:

- If operands are listed for a command, you must specify one. (If an operand is optional, the conditions under which it is optional are explained.)
- Optional parameters for a command are shown in brackets ([]).
- Variables are shown in italics (such as *nnn* for numbers).

Organization of this chapter and your needs

Review the chart for the types of commands covered in this chapter.

Types of Commands Covered in the Chapter	Commands Available	Page Number of Resource
Commands to access objects in batch	Primary Objects	p.543
	Authorizations	p.544
Online commands with additional functions in batch	END	p.552
	PRNT	p.569
Online commands with parameters or additional information	AMS	p.545
	BIND	p.547
	DISP	p.548
	OUTC	p.559
	START	p.571
	UTIL	p.572
Commands to control execution of statements	EXEC	p.554
	NOEXEC	p.558
	PDSEXEC	p.566
Commands to manage the PDS members with batch	PDSCOPY	p.562
	PDSEEDIT	p.564
	PDSLST	p.567
Commands for additional information	DUMP	p.549
	FLIST	p.555

General Commands for Specifying Objects in Batch

Overview

This unit covers the commands you must use to specify the type of object for processing.

Background about equivalent functions available online

Commands in this unit are equivalent to options to display objects for the !DB/WORKBENCH Primary Menu. When you use !DB/WORKBENCH online to display or perform operations on objects, you first use options on the Primary Menu to specify the type of object you want. For example, with online processing, you use the D option on the menu to display all databases.

Commands that specify objects

Review the chart to select the batch command that specifies the object you want to use. For example, if you want to act on databases, you use the DB command in the batch job.

Type of Object You Want to Use	Command to Use in Batch
All objects	AL
Aliases	AS
Databases	DB
Indexes	IX
Packages	PK
Plans	PL
Storage groups	SG
Table spaces	TS
Tables	TB
Views	VW

General Commands for Specifying Authorizations in Batch

Overview

This unit presents the different commands to specify authorizations for processing in batch.

Background about equivalent functions available online

When you use !DB/WORKBENCH online to display or perform operations on authorizations, you first use the Z option on the !DB/WORKBENCH Primary Menu to access the Authorizations Menu. You then specify the type of authorizations you want. For example, you use the D option on the Authorizations Menu to display authorizations for all databases.

Commands to specify authorizations

Review the chart to select the batch command for the authorization you want. You must use the required space between the command and its operand.

Authorizations	Command to Use in the Batch Job
Buffer pools	AU B
Collections	AU L
Columns for tables	AU C
Databases	AU D
Grantees	AU E
Packages	AU K
Plans	AU P
Storage groups	AU G
Table spaces	AU S
Tables	AU T
USERIDs	AU U
Views	AU V

AMS Command and Parameters for Batch

Description

The AMS command performs the same function in batch that it performs online. It generates Access Method Services LISTCAT, DEFINE, or DELETE statements for selected objects and their dependents. This command manages databases and explicitly defined table spaces or index spaces that DB2 has not allocated to a storage group.

Requirements for batch processing

AMS information is output on SYSPRINT. These operands and parameters are used with the AMS command in batch.

Operands and Parameters	Additional Information
LISTCAT	Lists the catalog entry for the dataset
DEFINE [ALL or ONLY] (Default=ALL)	Generates AMS DEFINE statements for all objects in the object list (ALL) or only for those table spaces or index spaces explicitly defined (ONLY)
DELETE [ALL or ONLY] (Default=ALL)	Generates AMS DELETE statements for all objects in the object list (ALL) or only for table spaces or index spaces that are explicitly defined (ONLY)

Requirements for using the LISTCAT operand with output directed to a file

You must include a LISTCPRT DD statement in the KTWBUTIL JCL. Omission of this DD statement will result in a corrupted SYSPRINT dataset if it is directed to a file.

Example of using the AMS command

In the example, you want to generate AMS statements to delete only table spaces that are explicitly defined (and not allocated by DB2 as a part of a storage group)

```
TS
NOFILT
TS=TS003
AMS DELETE ONLY
END
END
```

BIND Command and Parameters for Batch

Description

In batch, as with online processing, use the BIND command to generate bind statements for plans or packages.

Requirements for batch processing

These operands and parameters are used with the BIND command in batch.

Operands and Parameters	Additional Information
[ADD]	Adds the plan in the DB2 catalog
[REPLACE] [RETAIN] (Default=REPLACE)	Replaces the plan in the DB2 catalog and (optionally with RETAIN) retains the previous authorization on the plan

Examples using the BIND command

In the following example, you want to generate BIND commands for packages beginning with the letters QRS. You want to replace the plan in the DB2 catalog and retain previous authorizations for the plan.

```
PK
NOFILT
PKG=QRS*
BIND REPLACE RETAIN
END
END
```

In the following example, you want to generate BIND statements for one package identified by the letters QRS. You want the job output to show the package, creator, and owner information for the one package. The BATJCL.CNTL dataset contains the BIND statements for the QRS package.

```
PK
NOEXEC
FILT
PKG=QRS
S B
PDSCOPY TO=userid.BATJCL.CNTL(BIND) MOD
PRNT
END
END
```

DISP Command and Parameters for Batch

Description

The DISP command generates DISPLAY DATABASE commands. Using DISP, you can control the number of lines of output for databases, table spaces, and restricted databases.

Requirements for batch processing

The output is written to SYSTSPRT. These operands and parameters are used with the DISP command in batch processing.

Operands and Parameters	Additional Information
[ACTIVE]	Produces output of information about databases having table spaces or index spaces allocated to applications.
[RESTRICT]	Produces output of information about restricted databases.
[USE or LOCKS]	Produces output of correlation-ID and connection-ID information (USE) or output of lock (concurrency) information (LOCK)
[LIMIT(<i>nnn</i>) or LIMIT(*)] (Default=50)	Produces output of unlimited lines (*) or limits the number of output lines to a specific number (<i>nnn</i>)

Example of using the DISP command in batch

In the example, you want to generate DISPLAY DATABASE commands for restricted databases only. You want LOCKS and unlimited lines.

```

DB
DISP RESTRICT LOCKS LIMIT(*)
END
END

```

DUMP Command and the Batch Utility

Description and sample

The DUMP command is available with batch utilities only. It produces a dump of a range of DB2 pages from a table space or an index space. The illustration shows a sample of output from the DUMP command. Repeating lines are not printed.

```

1 DATASET DSNH.DSNDBD.DSQDBCTL.DSQTSTCL.I0001.A001, DB2 PAGE X'000000' (0)
00000000 00000000 0941D701 00000018 01000002 00000001 00180000 0000...
00000020 00000000 00000000 00000000 00000000 00000000 00000000 0000...
OFFSETS 00000040 THROUGH 00000FC0 SAME AS ABOVE 3
00000FE0 00000000 00000000 00000000 00000000 00000000 00000000 0000... 4

DATASET DSNH.DSNDBD.DSQDBCTL.DSQTSTCL.I0001.A001, DB2 PAGE X'000001' (1)
00000000 10000000 2A0BF201 00000110 00002A08 00000005 00000000 0000 ...
00000020 00000000 00000000 00000000 00000000 00000000 00000000 0000
OFFSETS 00000040 THROUGH 00000A60 SAME AS ABOVE
00000A80 00000000 00000000 00000000 00000000 00000000 00000000 0000 ...
00000AA0 00000000 00000000 00000000 00000000 00000000 00000000 0000 ...
OFFSETS 00000AC0 THROUGH 00000FC0 SAME AS ABOVE
00000FE0 00000000 00000000 00000000 00000000 00000000 00000000 0000 ...

DATASET DSNH.DSNDBD.DSQDBCTL.DSQTSTCL.I0001.A001, DB2 PAGE X'000002' (2)
00000000 10000000 29C26A01 00000202 06B00BF2 094C0100 00093800 0301 ...
00000020 40400006 D8E4C5D9 E8F18001 E2D8D340 40D5D640 40404040 4040 ...

```

- 1** Name of the dumped DB2 dataset
- 2** DB2 page number in decimal and hexadecimal. (Note that DB2 counts pages relative to zero; the first page is page number 0.)
- 3** Hexadecimal offset into each DB2 page
- 4** Page data in both hexadecimal and characters

Prerequisite for using the DUMP command

In order to access the table spaces or index spaces to be dumped, you must be appropriately authorized for access by the external security program in use at your site.

Format for using the DUMP command

This format is required to use the DUMP command.

```
DUMP DB(name)  
      IX(name) | TS(name)  
      FROMPAGE(n)  
      TOPAGE(n)
```

Background about requirements for FROMPAGE and TOPAGE

The value for *n* for FROMPAGE(*n*) and TOPAGE(*n*) is the value for the DB2 page number. (The first page of the data space is always page 0.) You can express the value as either hexadecimal or decimal. Review the chart for requirements for expressing the values.

Value You Want to Use	Requirement
Decimal	1-8 decimal digits
Hexadecimal	1-6 digits in the format X'nnn' (You must prefix the digits with an X and you must place the digits in quotation marks. Provide leading zeros.)

Requirements to use the DUMP command

The output is written to one of the reports provided by the batch utility. (For more information, see “Reports Produced by the Batch Utility” on page 495 in the chapter “Introduction to !DB/WORKBENCH JCL and Batch.”) These operands and parameters are used with the DUMP command available with the batch utilities.

Operands and Parameters	Additional Information
DB(<i>name</i>)	Names the database that contains the table space or index space that is to be dumped. (This operand is optional <i>if</i> the table space or index space name is unique.)
TS(<i>name</i>) IX(<i>name</i>)	Names the table space or index space to be dumped. (Specify the name without the leading qualifier, which is the creator ID.)
FROMPAGE(<i>n</i>) or FP(<i>n</i>)	Indicates the starting page number of the pages to be dumped.
TOPAGE(<i>n</i>) or TP(<i>n</i>)	Indicates the ending page number of the pages to be dumped.

Example of using the DUMP command

In the example, you want to produce a dump of DB2 pages for a table space named TS003 (a name that is unique) from the first page through page 3; you are using decimals to express the value for *n*.

```
DUMP TS(TS003)
      FROMPAGE(0)
      TOPAGE(3)
```

END Command and Options in Batch

Description

The END command serves the same purpose in batch as it serves online; however, it performs two additional functions in batch:

- Indicates the end of the commands that apply
- Terminates the batch utility

Background about equivalent function available online

When used online, the END command returns to the prior panel. For example, you could use the END command in these ways online:

- To return to the original object list panel after using a command that displays data about an object
- To return to the !DB/WORKBENCH Primary Menu because you have completed all tasks you want to perform with the type of object

Guidelines about using commands requiring an END command

When you use commands online, some commands are followed by the END command. For example, if you use the DEP command on the Plans panel, you must use the END command to redisplay the Plans panel. In general, commands that require the END command display details or object list panels with additional information. Requirements for using the END command with these commands varies for online and batch processing.

Type of Processing You Are Using	Requirement for the END Command	Example of Using Multiple Commands
Using a command that requires the END command in online processing	Use the END command after each of these commands. (You can, however, do other actions, such as printing, before using the END command.)	Your actions online: Use DEP Use PRNT Use END Use DBR Use PRNT Use END
Using a command that requires the END command in batch processing	Use the END command one time, no matter how many of this type of command you have in your job.	Your batch commands: DEP PRNT DBR PRNT END

Guidelines about concluding commands for an object

In batch, you use the END command after the last command you want to apply to an object. For example, if you want to print the tables and generate CREATE statements for the tables beginning with the letters TBTEST*, you use this batch job:

```
TB
TB=TBTEST*
PRNT
OUT
END
```

Guidelines about terminating the batch utility

In batch, you use the END command as the last command in the batch job. The END command serves the purpose of terminating the batch utility.

Examples of using the END command

The illustration shows two examples of batch jobs and the purposes of the END command in each job.

Example 1	Example 2
PL NOFILT PL=ABCD* DEP PRNT END 1 END 2 END 3	PL NOFILT PL=ABCD* DEP PRNT DBR PRNT END 1 END 2 END 3

- 1** Concludes one or more commands that require an END command (Equivalent to the action online of using the END command on the Dependents for All Plans panel to return to the Plans panel itself)
- 2** Indicates that no other command in the batch job applies to this type of object (Equivalent to the action online of using the END command to return to the !DB/WORKBENCH Primary Menu)
- 3** Terminates the batch utility

EXEC Command

Description

The EXEC command is available with the batch utility. It enables automatic execution or job submission. EXEC is the default mode for the batch utility. If you have used the NOEXEC command to save generated statements to the !DB/WORKBENCH user PDS, the EXEC command turns off the NOEXEC command.

Background about related commands

The NOEXEC command turns off the EXEC command. For information, see “NOEXEC Command and the Batch Utility” on page 558. Also, see “PDSEXEC Command and the Batch Utility” on page 566.

Background about the automatic submission and the defaults

The EXEC command is the default used for batch processing. Unless you use the NOEXEC command in a batch job, the batch utility for !DB/WORKBENCH automatically submits the job or executes the statements you have generated. These examples are typical.

Example of Contents of Your Batch Job	Action by the System
UTIL command	Generates a job stream; submits the job
OUT command	Generates CREATE statements; executes the DDL statements
BIND command	Generates BIND statements; executes the DSN statements

Example of using the EXEC command

The EXEC command is used in conjunction with the NOEXEC command. For an example of using the two commands in a batch job, see “NOEXEC Command and the Batch Utility” on page 558.

FLIST Command

Description

The FLIST command is available with the batch utility. It produces a formatted list of DB2 objects. You can specify an exact database name or you can use an asterisk either to represent a single character or at the end of a name to represent multiple characters.

Requirements to use the FLIST command

The output is written to one of the reports produced by the batch utility. (For more information, see “Reports Produced by the Batch Utility” on page 495 in the chapter “Introduction to !DB/WORKBENCH JCL and Batch.”) This parameter is optional with the FLIST command.

Operands and Parameters	Additional Information
[<i>dbname</i>]	Identifies the database to be listed (Omitting the name prints all objects authorized to the job name or USERID.)

Guidelines for controlling the objects on the list

You have three choices for controlling the objects on the list.

Quantity of Objects You Want	Method to Use
All objects authorized to the job name or your userid	FLIST
All objects associated with a specific database	FLIST <i>dbname</i>
All objects associated with one or more databases with common characters in the database name	FLIST followed by a name using asterisks (such as DBT00* or DB*006)

Guidelines for using an asterisk

You can use an asterisk (*) when specifying a dbname for the FLIST command. You use the asterisk in the same way you would if you were using it in a filter. If you need detailed information about using the asterisk, see the chapter “Filtering Object Lists” in *!DB/WORKBENCH Basic User's Guide*.

Items not shown in the example of output

Because the actual output of the FLIST command is 120 bytes wide, the illustration is scaled to fit the page and these items are removed from the sample:

- BUFFERPOOL (if any) assigned to this table space
- EDITPROC (Name for the edit procedure, if any, for the table)
- VALIDNPROC (Name for the validation procedure, if any, for the table)

Example of output from the FLIST command

In the example, you use the command **FLIST DSN8D2AP** to produce the output shown in the illustration.

		1 ----- DB/WORKBENCH DB2=DSNX -----		PAGE 7	
2	TS CREATOR	8	9	10	11
TBLSPACE	TABLE - INDEX NAME	IX SPACE	CATALOG	STORAGE GROUP	STATS? TS TB IX

DSN8S2DP	MYSYSADM 3 DSN82.TDEPT 4 SYN=MYSYSADM.TDEPT 5 IX=DSN82.XDEPT1 } IX=DSN82.XDEPT2 6 } IX=DSN82.XDEPT3 } VIEW=DSN82.VDEPT 7		DSNCAT	DSN8G20	YS
		XDEPT1	DSNCAT	DSN8G20	YS
		XDEPT2	DSNCAT	DSN8G20	YS
		XDEPT3	DSNCAT	DSN8G20	YS
TACTYPE	DSN82 DSN82.TACTYPE SYN=MYSYSADM.TACTYPE IX=DSN82.XACTYPE1 IX=DSN82.XACTYPE2 VIEW=DSN82.VACTYPE VIEW=DSN82.VSTAFAC1		DSNCAT	DSN8G20	YS
		XACTYPE1	DSNCAT	DSN8G20	YS
		XACTYPE2	DSNCAT	DSN8G20	YS
TEMPRAC	DSN82 DSN82.TEMPRAC SYN=MYSYSADM.TEMPRAC IX=DSN82.XEMPRAC1 IX=DSN82.XEMPRAC2 VIEW=DSN82.VEMPRAC		DSNCAT	DSN8G20	YS
		XEMPRAC1	DSNCAT	DSN8G20	YS
		XEMPRAC2	DSNCAT	DSN8G20	YS

FLIST report fields

The fields at the top of the report specify the following information.

- 1** Database name for the database
- 2** Table space name for the table space
- 3** Creator ID for the table space
- 4** Table name in *creator_ID.table_name* format
- 5** Synonyms defined for this table (if any) in *SYN=creator_ID.synonym_name* format
- 6** Index name in *IX=creator_ID.index_name* format
- 7** View name in *VIEW=creator_ID.view_name* format
- 8** Name of the index space that contains the index
- 9** Name of VSAM intersystem communications function (ICF) catalog (or its alias) that owns this table space or index space
- 10** Storage group, if any, for this table space or index space
- 11** Availability of statistics for the table space, table, or index (A YS in the column indicates the DB2 RUNSTATS utility has been run for the object.)

NOEXEC Command and the Batch Utility

Description

The NOEXEC command is available with the batch utilities. It disables automatic execution or job submission. NOEXEC permits you to generate DDL statements and save these utility job streams in the !DB/WORKBENCH user PDS for later submittal. You can place this control statement anywhere in the input stream and use it as often as needed.

Background about related commands

The EXEC command turns off the NOEXEC command. You can also use the PDSEXEC command to execute a PDS member you just created.

Examples of using the NOEXEC command

The chart shows two examples of using the NOEXEC command.

Way You Want to Use the NOEXEC Command	Batch Job to Use
Deferring execution and copying control statements to a PDS member for later execution.	<pre> TS NOFILT NOEXEC TS=TS006 UTIL COPY PDSCOPY, TO=CANDLE.SCHED.TEST(JOB02) END </pre>
Combining statements from multiple utilities	<pre> TS NOFILT NOEXEC TS=TS006 UTIL COPY UTIL REORG MOD UTIL COPY MOD EXEC UTIL MODIFY MOD END </pre>

OUTC Command and Parameters in Batch

Description

The OUTC command generates CREATE statements for all objects listed and for the dependent objects according to the parameters you use on the OUTC command.

Requirements for batch processing

These parameters are available with the OUTC command in batch.

Operands and Parameters	Statements Generated
ALL	CREATE statements for all objects
AS	CREATE statements for all aliases
IX	CREATE statements for all indexes
RF	Statements for referential alters
ST	CREATE statements for all synonyms on tables (You must specify the table to use this parameter.)
SV	CREATE statements for all synonyms on views (You must specify the view to use this parameter.)
TB	CREATE statements for all tables
TS	CREATE statements for all table spaces
VW	CREATE statements for all views

Background about equivalent functions online

When used online, the OUTC command displays the OutDDL options panel. This panel has options to specify which type of dependent objects you want, ranging from table spaces to aliases; however, the panel does not include an option to specify databases. Because the database option does not exist on the OutDDL Options panel, the OUTC command used on the Databases panel always results in CREATE statements for the databases on the object list.

Guidelines for generating the statements you want

Follow these guidelines to generate the statements you want.

Statements You Want to Generate	Guideline
Statements for databases	Be sure to specify DB as the type of object you want acted on by the batch job. (You do not specify a parameter with the OUTC command.)
Statements for any other supported object	Specify the parameter for the type of object with the OUTC command.

Example of using the OUTC command

The chart shows examples of using the OUTC command with different objects and dependents.

Statements You Want to Generate	Batch Job to Use
CREATE statements for all objects for the database DBDEA210	DB NOFILT DB=DBDEA210 OUTC ALL END END
CREATE statements for databases, table spaces, tables, and indexes for the database DBDEA210	DB NOFILT DB=DBDEA210 OUTC TS TB IX END END
CREATE statements for table space, table, and index for table space TSPEA210	TS NOFILT TS=TSPEA210 OUTC TS TB IX END END
ALTER statements for referential relationships	TB NOFILT TB=EMPLOY OUTC RF END END
CREATE statements for tables and synonyms for tables	TB NOFILT TB=EMPLOY OUTC TB ST END END

PDSCOPY Command and the Batch Utilities

Description

The PDSCOPY command is available with the batch utility. It copies the generated JCL to a destination you specify. For example, you could use the PDSCOPY command if you want to use a job-scheduling package instead of submitting the jobs as they are generated.

Prerequisite for using the PDSCOPY command

You can use the PDSCOPY command after the batch facility has generated a member in the !DB/WORKBENCH user PDS.

Format for using the PDSCOPY command

This format is required to use the PDSCOPY command.

```
PDSCOPY FROM=data_set_name(mbr name),  
TO=data_set_name(member name)
```

Background about requirements for the datasets

The default dataset for both the FROM and the TO parameters is the !DB/WORKBENCH user PDS. You do not need to specify the name of the dataset if it is the default. If you are copying from or to another dataset, you must provide a fully qualified dataset name without quotes.

If a command is too long for a command line, you indicate that the line continues by using a comma in one or more of these positions:

- Before the word *FROM*
- Before the word *TO*

Requirements for using the PDSCOPY command

These operands and parameters are used with the PDSCOPY command available with the batch utility.

Operands and Parameters	Additional Information
FROM=[data_set_name](member_name)	Indicates the member name to be copied (If you omit the FROM parameter, the newly created member becomes the default.)
TO=[data_set_name](member_name)	Indicates the member name for the copy
[MOD]	Adds the member to the end of the original member rather than replacing it

The PDSCOPY command is *not* designed to work inside either an edit macro or a !DB/WORKBENCH \$MACRO.

Examples of the PDSCOPY command

The chart shows examples of how to use the PDSCOPY command. (The examples also show where to position commas if the command is too long to fit on one line.)

Task You Want to Perform	Command to Use
Copy member USERIDA from the !DB/WORKBENCH user PDS to member UTILGEN1 in the !DB/WORKBENCH user PDS	PDSCOPY, FROM=(USERIDA), TO=(UTILGEN1)
Copy the member that was just generated into a different dataset	PDSCOPY, TO=TSTDS.TMP.PDS(CPYMBR)
Copy the member that was just generated into a different dataset and add its contents to the end of the original member	PDSCOPY, TO=TSTDS.TMP.PDS(CPYMBR) MOD

PDSEDIT Command and the Batch Utility

Description

The PDSEDIT command is available with the batch utilities. It invokes the ISPF Editor to use an ISPF edit macro with a specified member.

Background about related commands

Also see “NOEXEC Command and the Batch Utility” on page 558, and “PDSCOPY Command and the Batch Utilities” on page 562 and “PDSEXEC Command and the Batch Utility” on page 566.

Background about requirements for the dataset

The default dataset is the !DB/WORKBENCH user PDS. You do not need to specify the name of the dataset if it is the default. If you do specify a different dataset name, you must provide a fully qualified dataset name without quotes.

Requirements for using the PDSEDIT command

These operands and parameters are available with the PDSEDIT command.

Operands and Parameters	Additional Information
<i>[data_set_name][(member_name)]</i>	Indicates the member to be edited. If you omit the member name, the newly-created member becomes the default.
MAC(macro_name)	Specifies the name of the edit macro.

The PDSEDIT command is *not* designed to work inside either an edit macro or a !DB/WORKBENCH \$MACRO.

Example of using the PDSEEDIT command

In the example, you want to use the edit macro *EDIT* to edit member *LG000206* in the user PDS.

```
DB
NOEXEC
PDSEEDIT (LG000206) MAC(EDIT)
PDSCOPY TO=PROD.JCL(MANDANT1)
PDSEEXEC
END
```

PDSEEXEC Command and the Batch Utility

Description

The PDSEEXEC command is available with the batch utility. It executes the member that was just created. You also use this command if you include the NOEXEC command in the batch job and follow it with the PDSEEDIT command.

Background about related commands

Also see “NOEXEC Command and the Batch Utility” on page 558, “PDSEEDIT Command and the Batch Utility” on page 564, and “PDSCOPY Command and the Batch Utilities” on page 562.

Example of using the PDSEEXEC command

This example shows how you can use the PDSEEXEC command to execute the member you just created.

```
DB
NOEXEC
PDSEEDIT (LG000206) MAC(EDIT)
PDSCOPY TO=PROD.JCL(MANDANT1)
PDSEEXEC
END
```

PDSLIST Command and the Batch Utility

Description

The PDSLIST command is available with the batch utility. It lists the contents of a PDS member.

Background about requirements for the dataset

The default dataset is the !DB/WORKBENCH user PDS. You do not need to specify the name of the dataset if it is the default. If you do specify a different dataset name, you must provide a fully qualified dataset name without quotes.

Requirements for using the PDSLIST command

These operands and parameters are available with the PDSLIST command.

Operands and Parameters	Additional Information
<i>[data_set_name][(member_name)]</i>	Indicate the member to be listed. If you omit the member name, the newly-created member becomes the default.

Additional requirements to use the PDSLIST command

The PDSLIST command must be followed by the PRNT command.

Examples for the PDSLIST command

The chart contains some examples of using the PDSLIST command.

Task You Want to Perform	Commands to Use
List the contents of member <i>mymember</i> in the dataset <i>mydata</i>	PDSLIST mydata(mymember) PRNT
List the contents of the job stream you have just created	PDSLIST PRNT

PRNT Command and Operation in Batch

Description

The PRNT command serves the same general purpose it serves online; however, it is an essential command in batch. You must include the PRNT command in your batch job to receive your output.

Background about equivalent functions online

When you display object lists online, you can print the list of objects to your system printer by using the PRNT command.

Requirements for seeing reports in batch

If you want to be able to see the object lists that result from a command, you must also use the PRNT command.

Determining if the PRNT command works with a command

The PRNT command in batch is used to print a list that may be longer than one panel. Most of the panels in !DB/WORKBENCH are lists and therefore will print when the PRNT command is issued. This is true in the case of the L (LOAD) select. If you issue the L select in batch with the PRNT command, you will receive a printed listing.

However, some panels do not contain very much data and they are traditional panels, with fields placed in specific areas on the display screen. Since the data on these panels is not an internal list, the PRNT command cannot print the data. This is true in the case of the N (dataset INFO) select. If you issue the N select in batch with the PRNT command, the N select works but the PRNT command fails because the data is not in the format the PRNT command expects.

Examples of using the PRNT command

The chart shows examples of using the PRNT command.

Report You Want	Batch Job to Use
Dependents and DBRMs for all plans and also the plans with no DBRMs	<pre> PL NOFILT DEP PRNT DBR PRNT DBRMS=0 PRNT END END </pre>
Storage groups and storage groups by volume	<pre> SG PRNT VOLS PRNT PRNT END END END </pre>
Table spaces, tables, and indexes associated with the database named DB001	<pre> AL NOFILT DB=DB001 PRNT END END </pre>
Table space sets (showing names of tables in the set) for the table named TB001	<pre> TB NOFILT TB=TB001 TSSET PRNT END END </pre>
Table space sets (showing quantities of tables and indexes) for the table named TB001	<pre> TB NOFILT TB=TB001 S F SET PRNT END END </pre>

START Command and Parameters in Batch

Description

The START command generates START DATABASE commands. The START command enables a program or utility to access the specified database. You can also use this command to reset these types of errors:

- Pending deferred restarts
- Write errors
- Read-only accesses
- Utility controls

Requirement for batch processing

The DB2 output is on SYSTSPRT. These operands and parameters are used with the START command in batch.

Operands and Parameters	Additional Information
ACCESS(FORCE)	Forces a start to override stopped conditions
ACCESS(RO)	Provides read-only access
ACCESS(RW)	Provides read-and-write access
ACCESS(UT)	Provides for DB2 utility processing

UTIL Command and Functions Supported in Batch

Description

The UTIL command performs the same functions in batch that it performs when used from the command line online. It allows you to run a specified DB2 utility or to use options available to manage utilities.

Operands and Parameters	Additional Information
CHECK DATA	Checks table spaces and resets pending status (Execute after CHECK INDEX.)
CHECK INDEX	Checks indexes for discrepancies (Execute before CHECK DATA.)
COPY	Performs full or incremental copies as specified
DIAGNOSE	Lists SYSUTIL rows and DB2 directory contents for a table space or index space
MERGECOPY	Merges incremental image copies with themselves or the previous full image copy
MODIFY	Reads and updates
QUIESCE	Establishes a point at which a set of table spaces are in a consistent state
RECOVER	Recovers table spaces and indexes
RECOVER INDEX	Recovers indexes alone
REORG	Reorganizes table spaces and indexes on physical storage
REORG INDEX	Reorganizes only indexes on physical storage
REPORT	Reports on referential data
RUNSTATS	Reads and updates or reads only descriptive statistics on table spaces and indexes
RUNSTATS INDEX	Reads and updates or reads only descriptive statistics on indexes only
STOSPACE	Reads and updates statistics for SPACE columns

Options available for managing the utilities

You can also use these online functions in batch.

- UTIL OPTS *profile_name*
- UTIL *utility_name* MOD

If you need more information about using these options, see “Running Utilities” on page 337.

Examples of using the UTIL command in batch

The chart shows examples of using the UTIL command.

Utility You Want to Use	Batch Job to Use
Use the REORG utility: <ul style="list-style-type: none"> ● For the table spaces associated with the database named DB001 ● Using the utility profile named TOOLS 	<pre> TS NOFILT DB=DB001 UTIL OPTS TOOLS UTIL REORG END </pre>

Examples of using the UTIL command in batch (continued)

The chart continues examples of using the UTIL command.

Utility You Want to Use	Batch Job to Use
<p>Use the COPY, REORG, and MODIFY utilities:</p> <ul style="list-style-type: none"> ● For the table space named TS006 ● Deferring execution of the utilities ● Appending the utility statements so that the utilities generate a job stream to copy all datasets for the table space, to reorganize the table space, to copy after the REORG, and then to execute the prior statements and clean up using the MODIFY utility. 	<pre> TS NOFILT TS=TS006 NOEXEC UTIL COPY UTIL REORG MOD UTIL COPY MOD EXEC UTIL MODIFY MOD END </pre>
<p>Use the COPY utility, but copy the output to another PDS member for deferred submission:</p> <ul style="list-style-type: none"> ● For the table space named TS006 ● Using the utility profile named CANDLE ● Copying the statements to the dataset CANDLE.SCHED.TEST for the PDS member JOB01 	<pre> TS NOFILT NOEXEC TS=TS006 UTIL OPTS CANDLE UTIL COPY PDSCOPY, TO=CANDLE.SCHED.TEST(JOB001) END </pre>

Appendix A. Candle Customer Support

Introduction

Candle Corporation offers a comprehensive maintenance and support plan to ensure you realize the greatest value possible from your Candle software investments. We have more than 200 technicians worldwide, committed to providing you with prompt resolutions to your support requests.

Customer Support hours of operation are from 5:30 A.M. to 5:00 P.M., Pacific Time. In the event of an after-hours or weekend emergency, Candle's computerized call management system ensures that a technician will return your call within one hour. For customers located outside of North America, after-hours and weekend support is provided by Candle Customer Support locations in the United States.

Electronic Support

Candle provides information and support services using

- Candle's home page at www.candle.com. You can use the Candle Web site to
 - open problem records
 - access maintenance information
 - order products or maintenance
 - access IBM compatibility information
 - download fix packs for distributed products
 - read news and alerts
 - scan a list of scheduled Candle education classes
- Candle Electronic Customer Support (CECS), an electronic customer support facility. You can access this facility through the IBM Global Network. You can use CECS to
 - open problem records
 - search our database for solutions to known problems
 - look for answers to commonly asked questions
 - read news and alerts
 - scan a list of scheduled Candle education classes

Both CECS and the Candle Web site are available 24 hours a day, 7 days per week.

Telephone Support

Our support network consists of product specialists who work with you to solve your problem.

Candle uses an online problem management system to log and track all support requests. Your request is immediately routed to the appropriate technical resource.

When you call to report a problem, please have the following information:

- your Candle personal ID (PID) number
- the release level of the Candle product
- the release level of IBM or other vendor software
- identifying information and dates of recently applied maintenance to your Candle product or IBM product
- a detailed description of the problem (including the error message) and the events preceding the problem
- a description of any unusual events that occurred before the problem

Customer Support Phone Numbers

	Telephone	Fax
North America	(800) 328-1811	
	(310) 535-3636	(310) 727-4204
Europe		
Belgium/Luxembourg	+32 (0) 3 270 95 60	+32 (0) 3 270 95 41
France	+33 (0) 1 53 61 60 60	+33 (0) 1 53 61 06 16
Germany/Switzerland/ Austria	+49 (0) 89 54 554 333	+49 (0) 89 54 554 170
Italy - Freephone	800 780992	
Netherlands	+31 (0) 30 600 35 50	+31 (0) 30 600 35 10
Scandinavia	+46 (0)8 444 5940	+46 (0)8 623 1855
U.K.	+44 (0)161 437 5224	+44 (0)161 437 5225
(Southern Europe, Middle East and South Africa Agents call U.K.)		
Asia Pacific - English Hub		+61 2 9954 1818
Australia	+61 2 8912 9898	
Hong Kong	800 908 457	
India	+61 2 8912 9898	
Indonesia	0018 03061 2061	
Malaysia	1800 803 459	
New Zealand	0800 449 596	
Philippines	1800 1612 0096	
Singapore	800 616 2075	
Thailand	0018 00612 1045	
Asia Pacific - Japanese Hub	+81 3 3595 7150	+81 3 3595 7110
Asia Pacific - Korean Hub	+82 2 552 8744	+82 2 552 8746
Asia Pacific - Mandarin Hub	+88 62 2739 3223	+88 62 2378 5993
Asia Pacific e-mail address: ap_support@candle.com		

When your local support office is unavailable, you can contact Candle's North America support center. If USADirect® service is available in your country, use the 800 telephone number. If USADirect service is not available, ask your international operator for assistance in calling Candle's local (310) number.

Incident Documentation

You may be asked to send incident documentation to the Candle Customer Support Center. On the outside of all packages you send, please write the incident number given to you by the Customer Support representative.

Send tapes containing the incident information to the following address, unless directed otherwise by your Customer Support representative:

Candle Customer Support
Candle Support Center, *Incident number*
201 North Douglas Street
E1 Segundo, CA 90245

Send all other relevant documentation, such as diskettes or paper documentation, to the address provided by your Customer Support representative.

Ensuring Your Satisfaction with Customer Support

Candle Customer Support is committed to achieving high customer satisfaction ratings in all areas. These include

- connecting you to a support representative promptly
- providing you with the appropriate fixes
- answering support questions
- filling your shipping orders
- supplying documentation

If you have a concern that has not been resolved to your satisfaction, you can open a complaint ticket. All tickets are logged and tracked to ensure responsiveness and closure. Using the ticket information, a manager will contact you promptly to resolve your problem.

Index

Special Characters

[special character 35

A

access

Start a Database

required values 64

access exit 167, 171

access from command line

utilities 355

access from utilities menu

utilities 355

accessing

menu for utility profiles 454

accessing object lists 502

Action

Recover Indoubt

required values 75

activity log

displaying 144, 186

manipulating data 145

Adobe portable document format 16

allocating datasets

options available for 446

allocating sort work datasets 435

allocation of sort work datasets

examples 437

Alter Bufferpool

required values 86

ALTER BUFFERPOOL DB2 command 57

Alter Group Bufferpool

required values 87

ALTER statements

for referential integrity 215

AMS command

as batch command 545

AMS DEFINE statements 328

Archive Log

authority 83

DB2

options 57

managing the Archive Log (DB2) 57

Archive Log (*continued*)

required values 83

ARCHIVE LOG DB2 command 57

ARCHIVE privilege

Archive Log 83

Display Archive 84

Set Archive 84

ARVP option 126

AR–Archive Log

managing the Archive Log (DB2) 57

asterisk in SQL statements 37

Authids trace field 79

authority

Alter Bufferool 85

Alter Group Bufferool 87

Archive Log 83

Cancel DDF Thread 73

DBADM

creating a table 268

creating a table space 251

for index space 288

DBCTRL

creating a table 268

creating a table space 251

for index space 288

DBMAINT

creating a table 268

creating a table space 251

Display Archive 84

Display Bufferool 88

Display Connections 74

Display Database Status 62

Display Group 92

Display Group Bufferool 89

Display Locations 93

Display RLIMIT 69

Display Utility Status 82

hierarchy 99

INDEX privilege

for index space 288

RECOVER BSDS 91

Recover Indoubt 75

Reset Generic LU 94

Reset Indoubt 76

SELECT 179

- authority (*continued*)
 - Set Archive 84
 - Start a Database 64
 - Start DDF 67
 - Start RLIMIT 70
 - Stop a Database 65
 - Stop a Utility 82
 - Stop DB2 66
 - Stop DDF 68
 - Stop RLIMIT 71
 - SYSADM 238
 - creating a table 268
 - creating a table space 251
 - creating storage groups 308
 - for index space 288
 - traces 78
- authorizations
 - batch commands for 544
 - changing catalog prefixes for IDF 163
 - changing configuration information 165
 - CREATESG privilege
 - creating storage groups 308
 - creating a database 238
 - creating a table 238
 - CREATETAB privilege for database 268
 - DBADM 268
 - DBCTRL 268
 - DBMAINT 268
 - SYSADM 268
 - creating a table space
 - CREATETS privilege 251
 - DBADM 251
 - DBCTRL 251
 - DBMAINT 251
 - SYSADM 251
 - differences between menu and select 103
 - display
 - from an object list panel 106
 - from the menu 105
 - quick access from the menu 105
 - displaying 102
 - displaying profile datasets 189
 - displaying the activity log 186
 - displaying the system PDS 180
 - displaying ZPARMs values 128
 - for accessing global print options 156
 - for accessing specify DB2 subsystem information 153, 159
 - for setting the profile datasets 191
 - Housekeeping 144
- authorizations (*continued*)
 - index space 288
 - restoring catalog copies 182
 - running a new extract 179
 - specifying in batch 544
 - SYSADM authority
 - creating storage groups 308
 - update 172
 - updating for users 145
 - using a different system PDS 184
 - using filters to control the display 107
 - using to manage the system 97–109
- Authorizations Menu
 - accessing 104
 - authority required to display 102
 - differences from the Auths (Z) select 103
 - displaying authorizations 105
- Auths (Z) 106
- Auths (Z) select
 - authority required to display 102
 - differences from the Authorizations Menu 103
- available default values 224
- avoiding duplicate names
 - dataset names 360
 - DDnames 360

B

- batch
 - accessing object lists 502
 - and available online functions 510
 - authorizations, commands to specify 544
 - command dictionary 539
 - command line functions 510
 - commands
 - for deferred execution 516
 - specifying authorizations 544
 - specifying objects 543
 - commands for objects 543
 - creating jobs 535
 - requirements for 499
 - filtering
 - methods 507
 - filtering examples 508
 - functions for utilities 510
 - jcl for 521
 - KTWBATCH 493
 - modifying DD statements 533, 537
 - for extract 537

- batch (*continued*)
 - modifying extract ID 534
 - modifying PDS names 532
 - multi-step job
 - example 515, 519
 - for deferred execution 519
 - numeric filtering 509
 - optional modifications 535
 - prerequisites for using 494
 - process for using 493
 - requirements for commands 512
 - S command 511
 - simple job
 - example 513, 514, 518
 - for deferred execution 518
 - for utilities 514
 - skeletons 493
 - background about 521
 - sorting in 511
 - specifying authorizations in 544
 - specifying objects 543
 - specifying system information 494
 - submitting job for execution 538
 - using commands 503
 - batch commands
 - AMS command 545
 - authorization selections 544
 - BIND command 547
 - DISP command 548
 - DUMP command 549
 - EXEC/NOEXEC commands 554, 558
 - FLIST command 555
 - OUTC command 559
 - PDSCOPY command 562
 - PDSEEDIT command 564
 - PDSEXEC command 566
 - PDSLIST command 567
 - PRNT command 569
 - START command 571
 - UTIL command 572
 - batch jobs
 - skeletons for running 524
 - batch processing
 - introduction 489
 - batch return codes 171
 - batch utilities
 - benefits of 490
 - END command 552
 - JCL for 489
 - prerequisites 494
 - batch utility
 - introduction 530
 - printing reports 495
 - reports produced by 495
 - batch utility commands
 - DUMP
 - DB operand 551
 - beginning a session
 - menu displayed 230
 - BIND command
 - as batch command 547
 - BookManager 18
 - BSDS
 - DB2
 - options for managing 58
 - BSDS privilege
 - RECOVER BSDS 91
 - Buffer Pool
 - default 224
 - buffer pool name
 - creating a table space 252
 - creating an index 289
 - buffer pools
 - available authorizations 103
- C**
- CANCEL command 275
 - Cancel DDF Thread
 - authority 73
 - managing threads (DB2) 55
 - CANCEL DDF THREAD DB2 command 55
 - Candle Electronic Customer Support (CECS) 575
 - CAPS command 30
 - CATALOG display field
 - FLIST command output panel 557
 - Catalog Prefix 161
 - catalog prefixes for IDF
 - authority 163
 - changing 163
 - required values 164
 - Catalog Prefixes for IDF panel
 - accessing the menu 164
 - Housekeeping options
 - Option 1.2 164
 - catalog table
 - changing values 145

Index

- catalog table prefixes for IDF
 - displaying 144
- changing device names 445
- CHECK DATA utility 367
- CHECK DATA utility with RI data 367
- CHECK INDEX utility 369
- Classes trace field 79
- clean up statistics in
 - STATDBAS table space 200
- CLIST library 152
- CLOSE
 - default 224
- CLOSE setting
 - creating a table space 252
 - creating an index 289
- clustering index 287
 - creating 290
- clustering index online
 - creating an index 289
- CMDS AFTER
 - values for 427
- CMDS BEFORE
 - values for 426
- CNTL dataset 152
- Collection ID Prefix 161
- column definition
 - adding comments 280
 - adding labels 280
 - prerequisites 278
 - process 279
 - values 278
- column positions in SQL statements 37
- Columns
 - available authorizations 103
 - defining in tables 272, 273
 - selecting to be indexed 298
- command dictionary 539
- commands
 - AMDS command
 - DEFINE statements 328
 - available in the ISPF Editor 275
 - CANCEL 275
 - DB2 53
 - Alter Bufferpool 57
 - Alter Group Bufferpool 57
 - ARCHIVE LOG 57
 - CANCEL DDF THREAD 55
 - DISPLAY ARCHIVE 57
 - Display Bufferpool 57
 - DISPLAY DATABASE 53
 - Display Group 58
- commands (*continued*)
 - DB2 (*continued*)
 - Display Group Bufferpool 57
 - Display Location 58
 - DISPLAY RLIMIT 54
 - DISPLAY THREAD 55
 - DISPLAY TRACE 56
 - DISPLAY UTILITY 56
 - for databases 61
 - MODIFY TRACE 56
 - RECOVER BSDS 58
 - RECOVER INDOUBT 55
 - Reset Generic LU 58
 - RESET INDOUBT 55
 - SET ARCHIVE 57
 - START DATABASE 53
 - START DDF 54
 - START RLIMIT 54
 - START TRACE 56
 - STOP DATABASE 53
 - STOP DB2 54
 - STOP DDF 54
 - STOP RLIMIT 54
 - STOP TRACE 56
 - TERM UTILITY 56
 - deferred execution 516
 - dictionary for batch 539
 - DUMP batch utility
 - DB operand 551
 - END 219
 - in batch 504
 - FILT 221
 - filtering
 - in batch 507
 - FLIST command, output from 556
 - for DB2 databases
 - options for selecting from lists 61
 - for tasks with authorizations 109
 - GEN 356
 - examples of use 357
 - in batch 539
 - LOCATE 220
 - NOFILT 221
 - in batch 507
 - OUT 214, 215
 - OUTC 214, 215
 - OUTX 233
 - PRNT
 - in batch 506
 - REBIND 391

- commands (*continued*)
 - REBIND PKGS 394
 - REBIND PLANS 397
 - requirements in batch jobs 512
 - RESET 298
 - REV 219
 - RFO 215
 - S in batch 511
 - SAVE 275
 - SQL
 - CAPS 30
 - KTCSQL 43
 - Q and QQ 43
 - ROLLBACK 36
 - use in batch 503
 - UTIL 355
- Commands menu, DB2
 - See* DB2 Commands menu
- commands, batch
 - deferred execution 516
 - specifying authorizations 544
 - specifying objects 543
- comments
 - adding to columns in a table 280
 - options on traces 78
 - using in SQL statements 37
- COMPRESS
 - default 224
- COMPRESS setting
 - creating a table space 252
 - estimating 119
- Compression dictionary
 - using 410
- compression routines effecting
 - System PDS 178
- configuration information
 - authority to change 165
 - how to change 165
- Connection ID
 - Display Connections
 - required values 74
- Connection name
 - Recover Indoubt
 - required values 75
- conventions 15
- Copy Data Sets (LOCAL) panel 438
- Copy Data Sets (RECOVERY) panel 440
- COPY utility 371
 - types of copies available 372
- copy, incremental
 - dataset sizing
 - for migrated objects 447
- Correlation ID
 - Recover Indoubt
 - required values 75
- create database
 - specifying a new storage group 244
 - specifying an existing storage group 242
 - using values from an existing database 240
- Create Database panel 237, 239
 - how to keep from creating database 237
- Create Index panel 292
 - accessing 286
- Create Menu 214, 215, 336
 - accessing the Create Database panel 237
 - accessing the Create Index panel 286, 292
 - accessing the Create Storage Group panel 307
 - accessing the Create Table Space panel 250
 - defining VSAM datasets 215
 - displayed at the beginning of a session 230
 - how to access 217
 - introduction 209
 - order of creation 218
 - selecting options 216
 - using defaults 223
- CREATE statements
 - generating 214
 - how stored 226
 - keying of values 215
 - providing values 220
 - storing or saving 214
- Create Storage Group panel 244
 - accessing 307
- Create Table panel
 - accessing 267
 - accessing the Tables option 269
 - column definition 272
 - functions 272
- Create Table Space panel 256
 - accessing 250
- CREATEDBA privilege 238
- CREATEDBC privilege 238
- CREATESG privilege
 - creating storage groups 308
- CREATETAB privilege for the database
 - creating a table 268
- CREATETS privilege
 - creating a table space 251

Index

- creating dependent objects 227
 - creating Help for
 - user variables 478
 - creating indexes
 - purpose 287
 - terminology 287
 - creating objects 214
 - creating user variables 467
 - creating utility profiles 456
 - creation order, objects
 - Creator ID
 - default 224
 - customer support 575
 - customizing
 - BLKSIZ 528
 - KTWBATCH 526, 531
 - KTWUTJCL 528
 - REGION 528
 - sizing parameters 528
 - skeletons 526
 - SPCCTL 528
 - customizing utility profiles 456
- ## D
- DASD requirements
 - estimating 116
 - the fudge factor 118
 - with data compression 119
 - dashes in SQL statements 37
 - data sharing 92
 - using this option 92
 - data sharing management
 - options 92
 - Database
 - available authorizations 103
 - create
 - specifying a new storage group 244
 - specifying an existing storage group 242
 - using values from an existing database 240
 - creating 216, 235–245
 - authorizations 238
 - default 224
 - displaying history 173
 - Grantor DBADM 107
 - how to create 239
 - how to keep from creating 237
 - prerequisites for creating 238
 - storage groups 238
 - database management
 - DB2 command options 53
 - display status 53
 - options 61
 - start a list of databases, table spaces, and indexes 53
 - stop a list of databases, table spaces, and indexes 53
 - database name
 - Start a Database
 - required values 64
 - Stop a Database
 - required values 65
 - dataset
 - global information 150
 - authority 150
 - displaying 150
 - Dataset Containing DSN1COPY 161
 - Dataset Containing DSNHDECP 161
 - Dataset Containing DSNTIAD 162
 - Dataset Containing DSNTIAUL 162
 - Dataset Containing DSNZPARMS 162
 - dataset sizing
 - for index recovery
 - when object is missing 448
 - incremental copies of migrated objects 447
 - options for allocating 446
 - dataset, global
 - displaying information 144
 - datasets
 - displaying history 173
 - profile
 - status 194
 - datasets, sort work
 - allocation of 437
 - date format 167
 - DA–Display Archive
 - managing the Archive Log (DB2) 57
 - DB (database) filter
 - using for authorizations for tables 107
 - using on the Authorizations panel 107
 - DB operand
 - DUMP command
 - batch utility 551
 - DB2 catalog
 - old copies
 - erasing 182
 - resurrecting 182
 - using current system PDS 182
 - using different system PDS 182

- DB2 catalog (*continued*)
 - restoring 145
 - running a new 145
- DB2 commands 53
 - for databases 61
- DB2 Commands menu 50
 - accessing options 52
 - options 59
- DB2 Data Set Allocation Menu 336
- DB2 Database for !DB/Tools 161
- DB2 DSNEXIT Dataset 161
- DB2 DSNLOAD Dataset 161
- DB2 Location Name 161
- DB2 macros
 - changing 126
 - displaying 126
- DB2 management
 - command options 54
 - stop the DB2 subsystem 54
- DB2 parameters
 - migrating using DSNZPARMs 136
 - printing 137
- DB2 RUNLIB Dataset 161
- DB2 subsystem
 - changing information 159
 - location name 161
 - required authority 159
 - required values 160–161, 162
 - specifying information for 145
- DB2 subsystem ID 161
- DB2 utilities
 - COPY
 - types of copies available 372
 - creating
 - avoiding the END command 358
 - reviewing entries for accuracy 358
- DB2 version 161
- DB2-managed dataset
 - completing 263
- DB2-managed VSAM dataset
 - preparing 302
- DBADM authority
 - creating a table 268
 - creating a table space 251
 - Display Database Status 62
 - for index space 288
 - Start a Database 64
 - Stop a Database 65
- DBCTRL authority
 - creating a table 268
 - creating a table space 251
 - Display Database Status 62
 - for index space 288
 - Start a Database 64
 - Stop a Database 65
- DBDNAME display field
 - FLIST command output panel 557
- DBMAINT authority
 - creating a table 268
 - creating a table space 251
 - Display Database Status 62
 - Start a Database 64
 - Stop a Database 65
- DBRM library 152
- !DB/Tools
 - global configuration information
 - required values 155
 - global dataset information 150
 - authority 150
 - changing 150
 - displaying 150
 - required values 152
 - global print options
 - profile data administration
 - displaying 148
 - profile datasets
- DB/Tools Database 161
- !DB/Tools DB2 Configuration Information panel
 - accessing the menu 160
 - Housekeeping options
 - Option 1.1 160
- !DB/Tools Global Dataset Names panel
 - accessing the menu 151
 - Housekeeping options
 - Option 1.0 151
- !DB/WORKBENCH
 - configuring 145
 - processing information 145
- !DB/WORKBENCH Configuration Information panel
 - accessing the menu 166
 - Housekeeping options
 - Option 1.5 166
- !DB/WORKBENCH Primary Menu
 - accessing the DB2 Commands Menu 50
- !DB/WORKBENCH Processing Information panel
 - accessing the menu 170
 - Housekeeping options
 - Option 1.4 170

Index

- DC–Display Connections (Thread)
 - managing threads (DB2) 55
- DD statements
 - and KTWBATCH 533
 - modifying for batch 533
 - modifying for extract 537
- DDF
 - options for managing 54
- DDF management
 - options 67
- DD–Display Database Status
 - managing database option (DB2) 53
- default panel 223
- default utility profile 452
- defaults
 - Buffer Pool 224
 - CLOSE 224
 - Creator ID 224
 - Database 224
 - ERASE 224
 - FREEPAGE 224
 - Gbpcache 224
 - LOCK 224
 - New Object Waiting 224
 - Password 224
 - PCTFREE 224
 - PRIQTY 224
 - Return Wait 224
 - SECQTY 224
 - Storage Group 224
 - timing of displays 225
 - UNIQUE 224
 - values 225
 - available 224
 - restoring 223
 - when creating objects 223
 - VSAM Catalog 224
 - when creating objects 223
- Defer
 - default 224
- Defer setting
 - creating an index 289
- DEFINE statement 328
- dependent objects 218
 - creating 227
- Detail?
 - Display Connections
 - required values 74
- device (esoteric) names 445
 - changing 445
 - recognized by !DB/WORKBENCH 445
- device names
 - changing 445
 - recognized by !DB/WORKBENCH 445
- DIAG ABENDS utility 375
- DIAG SYSUTIL utility 377
- dictionary, command 539
- DISASTER REC utility 379
- DISP command
 - as batch command 548
- Display Active Traces
 - managing traces (DB2) 56
- Display Activity Log panel
 - accessing the menu 187
 - Housekeeping options 187
 - Option 7 187
- display after user variables modified
 - KTWPUTUS panel 474
- Display Archive
 - DB2
 - options 57
 - managing the Archive Log (DB2) 57
 - required values 84
- DISPLAY ARCHIVE DB2 command 57
- Display Arvhive
 - required values 84
- DISPLAY BUFFERPOOL DB2 command 57
- Display Connections
 - authority 74
 - required values 74
- Display Connections (Thread)
 - managing threads (DB2) 55
- DISPLAY DATABASE DB2 command 53
- Display Database Status
 - authority 62
 - managing database option (DB2) 53
- display fields
 - CATALOG
 - FLIST command output panel 557
 - DBDNAME
 - FLIST command output panel 557
 - INDEX NAME
 - FLIST command output panel 557
 - IX SPACE
 - FLIST command output panel 557
 - STATS TS TB IX
 - FLIST command output panel 557
 - STORAGE GROUP
 - FLIST command output panel 557

- display fields (*continued*)
 - TABLE
 - FLIST command output panel 557
 - TBLSPACE
 - FLIST command output panel 557
 - TS CREATOR
 - FLIST command output panel 557
 - VIEWNAME
 - FLIST command output panel 557
- Display Group
 - authority 92
 - required values 92
- DISPLAY GROUP BUFFERPOOL DB2
 - command 57
 - ALTER GROUP BUFFERPOOL DB2
 - command 57
- display irregularities 178
- Display Locations
 - authority 93
 - managing threads (DB2) 55
 - required values 93
- DISPLAY privilege
 - Display Archive 84
 - Display Bufferpool 88
 - Display Connections 74
 - Display Database Status 62
 - Display Group 92
 - Display Group Bufferpool 89
 - Display Locations 93
 - Reset Generic LU 94
- Display RLIMIT
 - authority 69
 - managing the RLF option (DB2) 54
 - required values 69
- DISPLAY RLIMIT DB2 command 54
- Display System PDS Members panel
 - accessing the menu 181
 - Housekeeping options 181
 - Option 5 181
- DISPLAY THREAD DB2 command 55
- DISPLAY TRACE DB2 command 56
- DISPLAY UTILITY DB2 command 56
- Display Utility Status
 - authority 82
 - managing utilities (DB2) 56
 - required values 82
- DISPLAYDB privilege
 - Display Database Status 62
- displaying profile datasets
 - authority 189
- displaying the activity log
 - authority 186
- displaying the system PDS
 - authority 180
- displays, timing of
 - defaults 225
- distributed data facility
 - See* DDF
- DO LATER option 229
 - implications for beginning another
 - session 229
- DO NOW option 233
- documentation set 17
- documentation, online 18
- drop statistics in
 - STATDBAS table space 200
- DR–Display RLIMIT
 - managing the RLF option (DB2) 54
- DS1COPY 159
- DSN1COMP utility 119
- DSN1COPY 161
- DSN6ARVP 126
- DSN6FAC 126
- DSN6GRP 126
- DSN6LOGP 126
- DSN6SPRM 126
- DSN6SYSP 126
- DSNHDECP option 126
- DSNTIAD Plan Name 161
- DSNTIAR 159, 161
- DSNTIAUL Plan Name 161
- DSNZPARAM 159, 161
 - available functions 126
 - using 123–137
- DSNZPARAM Member Name 161
- DSNZPARAMs facility
 - available functions 126
 - changing information
 - installation 126
 - restart 126
 - displaying information
 - displaying parameters 131
 - editing macros 134
 - editing parameters 131
 - generating macros 127, 134
 - MACROS option 127
 - MIGRATE option 127
 - migrating DB2 parameters 136

Index

- DSNZPARMs facility (*continued*)
 - migrating values 127
 - performing system tasks 127
 - PRINT REPT option 127
 - printing DB2 parameters 137
 - reporting 127
 - writing to a JES printer 127
 - DSNZPARMs menu
 - accessing 128
 - quick access 129
 - options 130
 - DSNZPARMs panels 131
 - DT–Display Active Traces
 - managing traces (DB2) 56
 - DUMP command
 - as batch command 549
 - DUMP command operands
 - batch utilities 551
 - duplicate names, avoiding
 - dataset names 360
 - DDnames 360
 - duplicate table names, avoiding 270
 - DU–Display Utility Status
 - managing utilities (DB2) 56
 - Dynamic SQL Facility
 - See* SQL
- ## E
- E (Editor) function 274
 - Editor, ISPF 274
 - commands 275
 - errors 275
 - navigation 275
 - operations 275
 - electronic customer support 575
 - END command 219
 - as batch command 552
 - in batch utility 552
 - ending a session
 - saving CREATE statements 228
 - ERASE
 - default 224
 - ERASE setting
 - creating a table space 252
 - VSAM dataset values for index space 289
 - errors
 - in the ISPF Editor 275
 - recovering from SQL 34
 - esoteric (device) names 445
 - esoteric names
 - changing 445
 - recognized by !DB/WORKBENCH 445
 - estimating
 - indexes
 - duplicate keys 119
 - key length 119
 - estimator
 - using to create objects 222
 - values
 - Priqty 263
 - estimators
 - accessing from the Primary Menu 114
 - accessing when creating objects 115
 - available functions 116
 - how they work 116
 - Index Space
 - values reported 120
 - Index Space Estimator 116
 - Table Space
 - values reported 120
 - Table Space Estimator 116
 - exclamation point
 - using in SQL 35
 - Exclude (X) select 298
 - EXEC/NOEXEC commands
 - as batch commands 554, 558
 - executing batch job 538
 - existing table space
 - specifying values from 257
 - exits
 - access 167
 - extract 167
 - options 167
 - R-O PDS 167
 - extract
 - catalog
 - resurrecting and erasing old copies 182
 - displaying historical information 145
 - number of 178
 - number per day 178
 - referential integrity information 145, 178
 - restoring a DB2 catalog 145
 - running a new 178
 - authority 179
 - running a new DB2 catalog 145
 - statistics collection during 201
 - tables in the STATDBAS table space 198
 - timestamp 178
 - effect on number per day 178

extract description 160
 extract exit 167
 extract history
 displaying 173
 Extract History panel
 accessing the menu 174
 Housekeeping options 174
 Option 3 174
 extract ID 160
 and KTWBATCH 534
 modifying for batch 534
 significance in batch 534
 extract PLAN name 166
 extract process
 STEP50
 clean up statistics 200
 drop statistics 200
 extract sort 171
 extract, running a
 gathering statistical data 179
 informational return code associated
 with 179

F

FAC option 126
 features
 changes 19
 new 19
 field
 job card
 specifying values for 429
 job statements
 specifying values for 429
 field descriptions
 traces 79
 FILT command 221
 filtering
 batch examples 508
 batch methods 507
 in batch 507
 NOT EQUAL filter in batch 508
 numeric examples 509
 statuses 509
 to solve problems 507
 use of NOFILT 507
 values 509

filters
 selections from lists 221
 using DB (database) on the Authorizations
 panel 107
 using for authorizations for tables 107
 using to control the display of
 authorizations 107
 FLIST command
 as batch command 555
 output from 556
 FREEpage 118
 default 224
 FREEpage value
 creating a table space 252
 VSAM dataset values for index space 289
 fudge factor 118
 functions
 E (Editor) 274
 SAVE 274

G

gathering statistical data
 while running a new extract 179
 informational return code associated
 with 179
 Gbpcache value
 creating a table space 252
 GEN command
 examples of use 357
 to bypass utility generation panels 356
 General Options panel 428
 job control 428
 QUIESCE utility statement
 value for the JOBCONTROL field 389
 utilities 428
 values for options 428
 generating unique names 430
 global configuration information
 required authority 153
 global print options
 required authority 156
 Grantees
 available authorizations 103
 Grantor DBADM database 107
 Grants 171
 GRP option 126

GTF trace field 79

H

Help template after modification
 user variables 485

Help template before modification
 user variables 484

Hiperspace 171

Housekeeping
 system information, specifying
 in batch 494

Housekeeping options 139
 accessing the menu 146
 accessing the menu options 147
 accessing the Primary Menu (quick
 access) 147
 authorization requirements 144
 considerations 143
 customizing your system 145
 displaying information 144
 Option 1 145, 148, 149
 profile data administration 148
 Profile Data Administration panel 149

Option 1.0
 !DB/Tools Global Dataset Names
 panel 151

Option 1.1
 !DB/Tools DB2 Configuration Information
 panel 160

Option 1.2
 Catalog Prefixes for IDF panel 164

Option 1.4
 !DB/WORKBENCH Processing Information
 panel 170

Option 1.5
 !DB/WORKBENCH Configuration
 Information panel 166

Option 2 172

Option 3 145, 174
 Extract History panel 174

Option 4 145, 178

Option 5 181
 Display System PDS Members panel 181

Option 6 145, 183
 Old (LAZARUS) Catalog Data panel 183

Option 7 145, 187
 Display Activity Log panel 187

using extracts 145

Housekeeping options (*continued*)
 when to use 143

I

ICCAT utility 382

ID, Connection
 Display Connections
 required values 74

ID, Correlation
 Recover Indoubt
 required values 75

ID, limit table
 Start RLIMIT
 required values 70

ID, Network
 Recover Indoubt
 required values 75

ID, Utility
 Display Utility Status
 required values 82
 Stop a Utility
 required values 82

IDF
 catalog prefixes
 authority 163
 changing 163
 required values 164
 displaying catalog table prefixes 144
 Interactive Display Facility 164

image copy
 dataset sizing
 for migrated objects 447
 types available 372

INDEX NAME display field
 FLIST command output panel 557

INDEX privilege
 for index space 288

Index Space Estimator 116
 values reported 120

index spaces
 authorizations 288
 creating 216, 283–303
 prerequisites 289
 datasets, DB2-created and managed 322
 estimating
 partitioned 291, 333
 different values 335
 setting up index key values 334

- index spaces (*continued*)
 - partitions 320
 - prerequisites 118–119
 - relationship to index 288
 - Space Estimator 322
 - indexes
 - clustering 287
 - creating 216, 283–303
 - clustering 290
 - partitioned 291
 - prerequisites 289
 - procedure 292
 - purpose 287
 - terminology 287
 - timing 288
 - unique 290
 - using default values 296
 - using values from an existing index 294
 - displaying history 173
 - maintenance 288
 - partitioned 287, 291
 - relationship to index space 288
 - types 287
 - unique 287
 - Indexes panel
 - CHECK INDEX utility 369
 - DIAG ABENDS utility 375
 - RECOVER INDEX utility 400
 - REORG INDEX utility 407
 - RUNSTATS utility 413
 - utilities 344
 - input panel, SQL
 - accessing 27
 - description 28
 - exiting 27
 - using uppercase and lowercase 30
 - INSERT method
 - for building tables 118
 - Installation Parameters
 - changing 126
 - displaying 126
 - Interactive Display Facility (IDF) 164
 - ISPF Editor 274
 - commands 275
 - errors 275
 - navigation 275
 - operations 275
 - ISPF MSGS dataset 152
 - ISPF PANELS dataset 152
 - ISPF profile
 - saving defaults 223
 - ISPF SKELS dataset 152
 - ISPF TABLES dataset 152
 - ISPF UTIL dataset 152
 - IX SPACE display field
 - FLIST command output panel 557
- J**
- jcl
 - for utilities and batch 521
 - JCL for utilities 489
 - job card 171
 - specifying values for 429
 - job control
 - incrementing the job name 429
 - options with utilities 428
 - values for options 428
 - job name
 - incrementing 429
 - job statements
 - specifying values for 429
 - job, batch
 - deferred execution 518, 519
 - multi-step example 515
 - simple example 513
 - for utilities 514
 - jobs, batch
 - prerequisites 494
- K**
- key length
 - when estimating indexes 119
 - keys
 - displaying history 173
 - keys, duplicate 119
 - KTCSQL SQL command 43
 - KTCUTJCL 452
 - KTWBATCH 493
 - customizing 531
 - modifying DD statements 533
 - modifying extract ID 534
 - modifying PDS names 532
 - relationship to KTWUTJCL 535

KTWPUOMU
 utility profile menu options 458
 KTWPUTUS panel
 display after user variables modified 474
 KTWPUTUS panel library member
 user variables template panel 469
 KTWUT\$IX
 index to skeletons 524
 KTWUTJCL
 customizing 528
 driver for utility generation skeletons 525
 relationship to KTWBATCH 535

L

L (LOAD) select
 PRNT command, batch 569
 labels
 adding to columns in a table 280
 Lazarus facility
See restoring a DB2 catalog
 Lazarus Retention 167
 STATDBAS table space 201
 left bracket
 using in SQL 35
 List option 349
 lists
 selections from 220
 LOAD library 152
 LOAD method
 for building tables 118
 LOCATE command 220
 Location
 Display Connections
 required values 74
 Location trace field 79
 LOCK
 default 224
 Lockout conditions
 concerning the PDS files 165
 LOCKsize setting
 creating a table space 252
 log PDS 167
 LOGONLY
 recovery using 404
 LOGP option 126

M

macros
 DB2, changing 126
 DB2, displaying 126
 generating from the DSNZPARMs
 facility 127
 generating using DSNZPARMs 134
 MACROS option 127
 managing bufferpools
 ALTER BUFFERPOOL DB2 command 57
 ALTER GROUP BUFFERPOOL DB2
 command 57
 DISPLAY BUFFERPOOL DB2 command 57
 DISPLAY GROUP BUFFERPOOL DB2
 Command 57
 managing the system
 using authorizations 97–109
 Max block size for VIO 155
 memory
 beginning the session 230
 creating dependent objects 227
 ending a session 228
 how CREATE statements are stored 226
 interpreting lists of selections 227
 menus
 Administrative Functions
 accessing 146, 149, 174, 181, 183, 187
 Authorizations
 accessing 104
 authority required to display 102
 differences from the uths (Z) select 103
 displaying authorizations 105
 Catalog Prefixes for IDF panel 164
 Create 214, 215
 accessing the Create Database panel 237
 accessing the Create Index panel 286, 292
 accessing the Create Storage Group
 panel 307
 accessing the Create Table Space
 panel 250
 defining VSAM datasets 215
 displayed at the beginning of a session 230
 how to access 217
 introduction 209
 order of creation 218
 selecting options 216
 using defaults 223
 Create Menu 336
 DB2 Commands 50, 59

- menus (*continued*)
 - DB2 Commands (*continued*)
 - accessing options 52
 - DB2 Data Set Allocation Menu 336
 - !DB/Tools DB2 Configuration Information panel 160
 - !DB/Tools Global Dataset Names panel 151
 - !DB/WORKBENCH Configuration Information panel 166
 - !DB/WORKBENCH Primary Menu
 - accessing the DB2 Commands Menu 50
 - !DB/WORKBENCH Processing Information panel 170
 - Display Activity Log panel
 - Display System PDS Members panel
 - DSNZPARMs
 - accessing 128
 - accessing (quick access) 129
 - options 130
 - Extract History panel
 - for utilities 348
 - Housekeeping options
 - accessing the options 147
 - Old (LAZARUS) Catalog Data panel
 - Primary
 - accessing Housekeeping options (quick access) 147
 - accessing the Authorizations Menu 104
 - accessing the estimators 114
 - accessing the Housekeeping options menu 146
 - Process Member
 - displayed at the beginning of a session 230
 - executing accumulated statements 233
 - integrating 231
 - options for processing statements 232
 - Profile Data Administration panel
 - MERGE COPY utility 384
 - MIGRATE option 127
 - migrated objects
 - dataset sizing for incremental copies 447
 - mode
 - Stop DDF
 - required values 68
 - modifications, optional
 - batch jobs 535
 - modifications, recommended
 - for batch 532, 533, 534
 - for extract 537
 - Modify Trace
 - managing traces (DB2) 56
 - MODIFY TRACE DB2 command 56
 - MODIFY utility 387
 - modifying utility profiles 456
 - MSGS dataset 152
 - MT–Modify Trace
 - managing traces (DB2) 56
 - multiple utility profiles 452
 - MVS ID 160
- N**
- N (dataset info) select
 - PRNT command, batch 569
 - names, unique
 - dataset names
 - avoiding 360
 - DDnames
 - avoiding 360
 - generating
 - multiple jobs for multiple activities 430
 - multiple jobs for single activity 430
 - navigation
 - in the ISPF Editor 275
 - Network ID
 - Recover Indoubt
 - required values 75
 - New Object Waiting
 - default 224
 - NOFILT command 221
 - non-partitioned dataset
 - completing 263
 - NOT EQUAL filter in batch 508
- O**
- O (Out DDL) select 214, 215
 - OBID value 281
 - object list panels
 - accessing the DB2 Commands Menu 50
 - display authorizations 106
 - displaying authorizations 102
 - object lists
 - access in batch 502
 - options for selecting from 77

- objects
 - batch commands for 543
 - creating 214
 - accessing the estimators 115
 - determining the order of creation 218
 - keying of values 215
 - creating dependent 227
 - dependent 218
 - order of creation 218
 - specifying in batch 543
- objects, migrated
 - dataset sizing for incremental copies 447
- Old (LAZARUS) Catalog Data panel
 - accessing the menu 183
 - Housekeeping options 183
 - Option 6 183
- online documentation 18
- online functions
 - availability in batch 510
- online Help
 - utilities 423
- operations
 - in the ISPF Editor 275
- OPn trace field 79
- Option 1
 - Housekeeping 145
- Option 3
 - Housekeeping 145
- Option 4
 - Housekeeping 145
- Option 6
 - Housekeeping 145
- Option 7
 - Housekeeping 145
- options
 - Create Menu 214, 215
 - DB2
 - AR–Archive Log 57
 - CT–Cancel DDF Thread 55
 - DA–Display Archive 57
 - DC–Display Connections (Thread) 55
 - DD–Display Database Status 53
 - DR–Display RLIMIT 54
 - DT–Display Active Traces 56
 - DU–Display Utility Status 56
 - managing BSDS recovery 58
 - managing bufferpools 57
 - managing data sharing 58
 - managing databases 53
 - managing DB2 54
 - managing the Archive Log 57
 - options (*continued*)
 - DB2 (*continued*)
 - managing the DDF 54
 - managing the RLF 54
 - managing threads 55
 - managing traces 56
 - MT–Modify Trace 56
 - PD–Stop a Database 53
 - PF–STOP DDF 54
 - PR–Stop RLIMIT 54
 - PT–Stop a Trace 56
 - PU–Stop a Utility 56
 - RB–Display Group 58
 - RB–Display Location 58
 - RB–Recover BSDS 58
 - RB–Reset Generic LU 58
 - RI–Recover Indoubt 55
 - RS–Reset Indoubt 55
 - SA–Set Archive 57
 - SD–Start a Database 53
 - SF–START DDF 54
 - SR–Start RLIMIT 54
 - ST–Start a Trace 56
 - utilities 56
 - DB2 Commands menu
 - how they work 59
 - interpreting output 60
 - output 59
 - DO LATER 229
 - implications for beginning another session 229
 - DSNZPARMs menu 130
 - for comments on the trace 78
 - for testing SQL statements 42
 - Housekeeping
 - See* Housekeeping options
 - Process Member menu
 - DO NOW 233
 - UPDATE NOW 233
 - stopping the current 219
 - options exit 167
 - OPX trace field 79
 - OUT command 214, 215
 - Out DDL (O) select 214, 215
 - OUTC command 214, 215
 - as batch command 559
 - OutDDL Options panel 215
 - output
 - generated by DB2 Commands menu options 60

- output from the DUMP command panel 549
- output from the FLIST command panel 556
- output, displaying SQL 41
 - differences 41
- OUTX command 233

- P**
- Package authority 171
- Packages
 - available authorizations 103
- panels
 - Create Database 237, 239
 - how to keep from creating database 237
 - Create Index 292
 - accessing 286
 - Create Storage Group 244
 - accessing 307
 - Create Table
 - accessing 267
 - Create Table Space 256
 - accessing 250
 - Create Tables
 - accessing the Tables option 269
 - column definition 272
 - functions 272
 - default 223
 - DSNZPARMs 131
 - object list
 - accessing the DB2 Commands Menu 50
 - displaying authorizations 102, 106
 - OutDDL Options 215
 - output from the DUMP command 549
 - output from the FLIST command 556
 - redisplay the previous 219
 - SQL input
 - accessing 27
 - description 28
 - exiting 27
 - using uppercase and lowercase 30
 - Summary Report from the PRNT command 496
 - table column
 - interpreting data 299
 - Table Columns 273
- PANELS dataset 155
- parallel I/O operations and degree
 - Degree 391, 394, 397
- parallel processing 428
- parameters
 - ROSHARE 240, 242, 244, 281
- partition independence
 - and processing 428
- partition value
 - creating a table space 252
- partitioned index 287, 291
 - partitioned 291
- partitioned index space 291, 333
 - setting up index key values 334
- partitioned table spaces 248
 - creating 254
- partitions
 - different values 335
- partitions and VSAM datasets 320
- Password
 - default 224
- passwords 167
- PCTFREE 118
 - default 224
- PCTfree value
 - creating a table space 252
 - VSAM dataset values for index space 289
- PDS
 - log 167
 - system 166
 - user 166
- PDS files
 - lockout conditions 165
- PDS member
 - changing names 38
- PDS names
 - and KTWBATCH 532
 - modifying for batch 532
- PDSCOPY command
 - as batch command 562
 - does not work in macros 563
- PDSEEDIT command
 - as batch command 564
 - does not work in macros 564
- PDSEEXEC command
 - as batch command 566
- PDSLIST command
 - as batch command 567
- PD-Stop a Database
 - managing database option (DB2) 53
- PF-STOP DDF
 - managing the DDF option (DB2) 54

Index

- PGSIZE 119
 - PLAN names
 - extract 166
 - SQL 166
 - Plans
 - available authorizations 103
 - Plans trace field 79
 - portable document format, Adobe 16
 - prerequisites for creating
 - user variables 468
 - prerequisites for creating Help
 - user variables 479
 - Primary Menu
 - accessing the Authorizations Menu 104
 - accessing the estimators 114
 - accessing the Housekeeping options menu 146
 - PRINT REPT option 127
 - PRIQTY
 - default 224
 - Priqty value 263
 - creating a table space 252
 - VSAM dataset values for index space 289
 - privileges
 - CREATEDBA 238
 - CREATEDBC 238
 - CREATESG
 - creating storage groups 308
 - CREATETAB (for the database)
 - creating a table 268
 - CREATETS
 - creating a table space 251
 - PRNT command
 - as batch command 569
 - PRNT command, batch
 - L (LOAD) select 569
 - N (dataset info) select 569
 - PRNT output 569
 - Problem Resolution
 - display irregularities 178
 - Process Member menu
 - displayed at the beginning of a session 230
 - DO NOW 233
 - executing accumulated statements 233
 - integrating 231
 - options for processing statements 232
 - UPDATE NOW 233
 - utilities 359
 - profile
 - data administration 148
 - displaying 148
 - datasets
 - global configuration
 - global datasets
 - required values 152
 - ISPF
 - saving defaults 223
 - Profile Data Administration panel
 - accessing the menu 149
 - Housekeeping options 149
 - Option 1 149
 - PROFILE datasets
 - status 194
 - Profile Information Panel
 - required authority 191
 - profile options
 - considerations 452
 - saving 462
 - profile, utility
 - creating 456
 - customization 456
 - default 452
 - location for 452
 - modifying 456
 - profile,utility
 - possibility of multiple 452
 - profiles, utility
 - options available 453
 - setting 453
 - providing values
 - for CREATE statements 220
- PR–Stop RLF
 - managing the DDF option (DB2) 54
- PT–Stop a Trace
 - managing traces (DB2) 56
- PU–Stop a Utility
 - managing utilities (DB2) 56
- ## Q
- Q and QQ SQL commands 43
 - quiesce
 - Archive Log
 - required values 83
 - Stop DB2
 - required values 66

QUIESCE utility 389
 value for the JOBCONTROL field 389

R

RACF qualifier 171
 RB–data sharing
 managing data sharing (DB2) 58
 RB–Recover BSDS
 managing BSDS (DB2) 58
 Real Catalog Prefix 161
 REBIND command 391
 REBIND PKGS command 394
 REBIND PLANS command 397
 Recover BSDS 91
 authority 91
 managing BSDS (DB2) 58
 required values 91
 using this option 91
 RECOVER BSDS DB2 command 58
 RECOVER INDEX utility 400
 Recover Indoubt
 authority 75
 managing threads (DB2) 55
 required values 75
 RECOVER INDOUBT DB2 command 55
 RECOVER utility for table spaces 402
 Recovery using LOGONLY 404
 redisplay the previous panel 219
 referential integrity
 generating statements 215
 utilities
 data for referential integrity 346
 reinstate statistics in
 STATDBAS table space 200
 REORG INDEX utility 407
 REORG utility for table spaces 409
 REPORT utility 411
 reports
 from the DSNZPARMs facility 127
 interpreting 40
 REPORTS DD statement
 batch utility 569
 SYSPRINT
 batch utility 569
 REPORTS DD statement report
 PRNT
 batch utility 569
 RESET command 298
 Reset Generic LU
 authority 94
 required values 94
 Reset Indoubt
 managing threads (DB2) 55
 required values 76
 RESET INDOUBT DB2 command 55
 resource limit facility
 See RLF management
 RESTART option 126
 Restart Parameters
 changing 126
 displaying 126
 restoring a DB2 catalog 145, 182
 restoring catalog copies
 authority 182
 restoring default values 223
 restrictions
 adding to tables 281
 Return Wait
 default 224
 REV command 219
 RFO command 215
 RI–Recover Indoubt
 managing threads (DB2) 55
 RLF management
 options 54, 69
 RLIMIT
 Display
 authority 69
 required values 69
 Start
 authority 70
 required values 70
 Stop
 authority 71
 required values 71
 RMIDs
 selecting 77
 RMIDs trace field 79
 R-O PDS exit 167
 ROLLBACK command 36
 ROSHARE parameter 240, 242, 244, 281
 running a new extract
 gathering statistical data 179
 informational return code associated
 with 179

RUNSTATS utility 413

S

SAVE command 275

SAVE function 274

saving CREATE statements

at the end of a session 228

SA–Set Archive

managing the Archive Log (DB2) 57

SD–Start a Database

managing database option (DB2) 53

SECQTY

default 224

Secqty value

creating a table space 252

VSAM dataset values for index space 289

security, maintaining 99

segmented table spaces 248

creating 253

SEGsize 119

SEGsize setting

creating a table space 252

SEGsize value

creating a segmented table space 253

SELECT authority 179

selecting from lists 220

using filters 221

selections

implication of memory on interpreting

selects

Auths (Z)

authority required to display 102

Exclude (X) 298

for displaying system PDS members 181

for displaying the activity log 188

for erasing old extract catalog copies 183

for resurrecting old extract catalog copies 183

for selections from lists 220

for tasks with authorizations 109

O (Out DDL) 214, 215

Out DDL (O) 214, 215

using with DB2 commands for databases 61

Util (U) 354

X (Exclude) 298

Z (Auths)

authority required to display 102

differences from the Authorizations

Menu 103

semicolon in SQL statements 37

Set Archive

DB2

options 57

managing the Archive Log (DB2) 57

required values 84

SET ARCHIVE DB2 command 57

SF–START DDF

managing the DDF option (DB2) 54

sizing

dataset 446

for incremental copies of migrated

objects 447

for index recovery

when object is missing 448

sizing parameters

BLKSIZ 528

customizing 528

REGION 528

SPCCTL 528

sizing, dataset

for index recovery

when object is missing 448

for utility generation 446

skeletons

background about 524

customizing 526

customizing KTWBATCH 531

customizing KTWUTJCL 528

driver for utility generation skeletons 525

for batch 524

for online 524

for online utilities 524

for running batch jobs 524

for utility generation 524

index to 524

KTWUT\$IX 524

KTWUTJCL 525

used in batch jobs 523

skeletons dataset 155

skeletons, batch

background about 521

SMF trace field 79

sort files

allocation

examples 437

sort work datasets

allocating 435

- sorting
 - in batch 511
- Space Estimator 322
 - Priqty value 263
 - using to create objects 222
- space name
 - Start a Database
 - required values 64
 - Stop a Database
 - required values 65
- ! special character 35
- special characters
 - using in SQL statements 37
- SPRM option 126
- SQL
 - accessing the input panel 27
 - commands
 - KTCSQL 43
 - Q and QQ 43
 - description of the input panel 28
 - displaying output 41
 - differences 41
 - entering statements directly 31
 - executing statements
 - entered directly 32
 - retrieved 33
 - exiting the input panel 27
 - interpreting reports 40
 - keyboard mapping 35
 - recovery from errors 34
 - statements
 - changing the PDS member 38
 - column positions 37
 - executing those in a different PDS 39
 - executing those stored in a PDS member 38
 - options for testing 42
 - procedure for editing and testing 43
 - storing in a PDS member 36
 - symbols required in a PDS member 37
 - testing single and multiple lines 43
 - testing within a program 42
 - using in !DB/WORKBENCH 23–43
 - using the exclamation point 35
 - using the left bracket 35
 - using uppercase and lowercase 30
- SQL PLAN name 166
- SQLID source 167
- SRV trace field 79
- SR–Start RLIMIT
 - managing the RLF option (DB2) 54
- Start a Database
 - authority 64
 - managing database option (DB2) 53
 - required values 64
- Start a Trace
 - managing traces (DB2) 56
- start accumulating statistics
 - in the STATDBAS table space 199
- START command
 - as batch command 571
- START DATABASE DB2 command 53
- START DDF
 - authority 67
 - managing the DDF option (DB2) 54
 - required values 67
- START DDF DB2 command 54
- Start RLIMIT
 - authority 70
 - managing the RLF option (DB2) 54
 - required values 70
- START RLIMIT DB2 command 54
- START TRACE DB2 command 56
- STARTDBD privilege
 - Start a Database 64
- STATDBAS table space
 - accumulating statistics 199
 - clean up statistics 200
 - controlling the size of tables 203
 - drop statistics 200
 - growth in size of tables 201
 - not required to run !DB/WORKBENCH 199
 - number of days of statistics to keep 199
 - reinstate statistics gathering 200
 - requirements at installation 199
 - stop accumulating statistics 199
 - tables required for statistics 199
- statements
 - accumulated, executing 233
- ALTER
 - for referential integrity 215
- CREATE
 - generating 214
 - how stored 226
 - keying of values 215
 - providing values 220
 - saving at the end of a session 228
 - storing or saving 214

Index

- statements (*continued*)
 - executing stored statements 232
 - interpreting reports 40
 - options for processing 232
 - processing from utilities 359
 - SQL
 - changing the PDS member 38
 - column positions 37
 - differences in displaying output 41
 - displaying output 41
 - entering directly 31
 - executing those entered directly 32
 - executing those in a different PDS 39
 - executing those retrieved 33
 - executing those stored in a PDS
 - member 38
 - options for testing 42
 - procedure for editing and testing 43
 - storing in a PDS member 35, 36
 - symbols required in a PDS member 37
 - testing single and multiple lines 43
 - testing within a program 42
- statistics
 - collection during extract 201
 - required tables in STATDBAS table space 199
- STATS TS TB IX display field
 - FLIST command output panel 557
- STATTIME column
 - in tables in STATDBAS table space 198
- statuses
 - filtering in batch 509
- STEP50
 - drop statistics
 - in the extract process 200
- Stop a Database
 - authority 65
 - managing database option (DB2) 53
 - required values 65
- Stop a Trace
 - managing traces (DB2) 56
- Stop a Utility
 - managing utilities (DB2) 56
 - required values 82
- Stop a Utility Status
 - authority 82
- stop accumulating statistics
 - STATDBAS table space 199
- STOP DATABASE DB2 command 53
- Stop DB2
 - authority 66
 - required values 66
- STOP DB2 command 54
- STOP DDF
 - authority 68
 - managing the DDF option (DB2) 54
 - required values 68
- STOP DDF DB2 command 54
- Stop RLF
 - managing the DDF option (DB2) 54
- Stop RLIMIT
 - authority 71
 - required values 71
- STOP RLIMIT DB2 command 54
- STOP TRACE DB2 command 56
- STOPALL privilege
 - Stop DB2 66
- STOPDB privilege
 - Stop a Database 65
- stopping the current option 219
- STORAGE GROUP display field
 - FLIST command output panel 557
- storage group name
 - creating a table space 252
 - VSAM dataset values for index space 289
 - VSAM datasets 330
- Storage Groups
 - available authorizations 103
 - creating 216, 305–313
 - authorizations 308
 - defining VOLSERS 308
 - modifying values from an existing storage group 312
 - prerequisites 308
 - procedure 309
 - terminology 308
 - using values from an existing storage group 310
 - default 224
 - prerequisites for creating a database 238
- Storage Groups panel
 - STOSPACE utility 417
 - utilities 344
- STOSPACE utility 417
- STUNITS 171
- ST–Start a Trace
 - managing traces (DB2) 56

- submitting batch job 538
- subpage setting
 - creating an index 289
- subsystem location name, DB2 161
- Summary Report from the PRNT command panel 496
- synonyms
 - using when creating tables 282
- SYSADM authority 238
 - Alter Bufferpool 85
 - Alter Group Bufferpool 87
 - Archive Log 83
 - Cancel DDF Thread 73
 - creating a table 268
 - creating a table space 251
 - creating storage groups 308
 - Display Archive 84
 - Display Bufferpool 88
 - Display Connections 74
 - Display Database Status 62
 - Display Group 92
 - Display Group Bufferpool 89
 - Display Locations 93
 - Display RLIMIT 69
 - for index space 288
 - RECOVER BSDS 91
 - Recover Indoubt 75
 - Reset Generic LU 94
 - Reset Indoubt 76
 - Set Archive 84
 - Start a Database 64
 - Start DDF 67
 - Start RLIMIT 70
 - Stop a Database 65
 - Stop a Utility 82
 - Stop DB2 66
 - Stop DDF 68
 - Stop RLIMIT 71
 - traces 78
- SYSCTRL authority
 - Alter Bufferpool 85
 - Alter Group Bufferpool 87
 - Archive Log 83
 - Cancel DDF Thread 73
 - Display Archive 84
 - Display Bufferpool 88
 - Display Connections 74
 - Display Database Status 62
 - Display Group 92
 - Display Group Bufferpool 89
 - Display Locations 93
 - Display RLIMIT 69
 - Recover Indoubt 75
 - Reset Generic LU 94
 - Reset Indoubt 76
 - Set Archive 84
 - Start DDF 67
 - Start RLIMIT 70
 - Stop a Utility 82
 - Stop DB2 66
 - Stop DDF 68
 - Stop RLIMIT 71
 - traces 78
- SYSCTRL authority (*continued*)
 - Display Locations 93
 - Display RLIMIT 69
 - RECOVER BSDS 91
 - Recover Indoubt 75
 - Reset Generic LU 94
 - Reset Indoubt 76
 - Set Archive 84
 - Start a Database 64
 - Start DDF 67
 - Start RLIMIT 70
 - Stop a Database 65
 - Stop a Utility 82
 - Stop DB2 66
 - Stop DDF 68
 - Stop RLIMIT 71
 - traces 78
- SYSINDEXES
 - column definitions 206
 - contents of 206
 - controlling the size of 203
- SYSOPR authority
 - Alter Bufferpool 85
 - Alter Group Bufferpool 87
 - Archive Log 83
 - Cancel DDF Thread 73
 - Display Archive 84
 - Display Bufferpool 88
 - Display Connections 74
 - Display Database Status 62
 - Display Group 92
 - Display Group Bufferpool 89
 - Display Locations 93
 - Display RLIMIT 69
 - Recover Indoubt 75
 - Reset Generic LU 94
 - Reset Indoubt 76
 - Set Archive 84
 - Start DDF 67
 - Start RLIMIT 70
 - Stop a Utility 82
 - Stop DB2 66
 - Stop DDF 68
 - Stop RLIMIT 71
 - traces 78
- SYSOUT Hold Class 155
- SYSP option 126
- SYSPRINT DD Statement 569

Index

- SYSPRINT report
 - PRNT
 - batch utility 569
- SYSTABLES
 - column definitions 205
 - contents of 205
 - controlling the size of 203
- SYSTABLESPACE
 - column definitions 204
- SYSTABLESPACES
 - conents of 204
 - controlling the size of 203
- system information
 - specifying, in batch 494
- system PDS 166
 - effect of compression routines 178
- system PDS members
 - displaying 180
- T**
- table
 - creating
 - authorizations 238
- table columns
 - selecting to be indexed 298
- Table Columns panel 273
- TABLE display field
 - FLIST command output panel 557
- Table Space Estimator 116
 - values reported 120
- Table Spaces
 - accessing the Create Table Space panel 250
 - available authorizations 103
 - building 118
 - create
 - partitioned 254
 - procedure 256
 - segmented 253
 - creating 216, 247–264
 - authorizations 251
 - prerequisites 252
 - values 252
 - datasets, DB2-created and managed 322
 - estimating
 - partitioned 248
 - different values 335
 - relationship to indexes 291
 - partitions 320
- Table Spaces (*continued*)
 - preparing VSAM datasets 248
 - prerequisites 118–119
 - RECOVER utility 402
 - REORG utility 409
 - segmented 248
 - Space Estimator 322
 - STATDBAS 196, 198
 - terminology 251
 - values
 - defaults 260
 - from an existing table space 257
- Table Spaces panel
 - CHECK DATA utility 367
 - CHECK DATA utility with RI data 367
 - CHECK INDEX utility 369
 - COPY utility 371
 - DIAG ABENDS utility 375
 - DIAG SYSUTIL utility 377
 - DISASTER REC utility 379
 - ICCAT utility 382
 - MERGE COPY utility 384
 - MODIFY utility 387
 - QUIESCE utility 389
 - REBIND command 391
 - REBIND PKGS command 394
 - REBIND PLANS command 397
 - RECOVER utility 402
 - REORG utility for table spaces 409
 - REPORT utility 411
 - RUNSTATS utility 413
 - utilities 344
- tables
 - adding restrictions 281
 - columns
 - selecting to be indexed 298
 - creating 216, 265–282
 - authorizations 268
 - avoiding duplicate names 270
 - column definition 272, 273
 - completing values 281
 - functions 272
 - synonyms 282
 - terminology 268
 - using the ISPF Editor 274
 - values 270
 - DB2 Catalog
 - displaying history 173
 - in the STATDBAS table space 198
 - partitioned
 - relationship to indexes 291

- tables (*continued*)
 - required to run statistics 199
 - SYSINDEXES 201
 - SYSTABLES 201
 - SYSTABLESPACE 201
 - using filters for authorizations 107
 - TABLES dataset 155
 - tables paces
 - displaying history 173
 - TBLSPACE display field
 - FLIST command output panel 557
 - telephone support 576
 - template after modification
 - user variables 477
 - template before modification
 - user variables 476
 - template for creating
 - user variables 470
 - Temporary Disk Symbolic Name 155
 - TERM UTILITY DB2 command 56
 - testing
 - SQL statements
 - multiple lines 43
 - procedure 43
 - single line 43
 - testing SQL statements 42
 - threads management
 - options 72
 - output 72
 - threads, DB2
 - options for managing 55
 - TIME
 - Archive Log
 - required values 83
 - timing of displays
 - defaults 225
 - Trace IDs trace field 79
 - TRACE privilege
 - traces 78
 - traces
 - authority 78
 - displaying
 - required values 81
 - displaying, monitoring, starting, and stopping
 - description of values 79
 - monitoring
 - required values 81
 - starting
 - required values 80
 - stopping
 - required values 81
 - traces management
 - options 77
 - options on menus 78
 - traces, DB2
 - options for managing 56
 - TS CREATOR display field
 - FLIST command output panel 557
 - Type
 - Display Connections
 - required values 74
 - types 131
- U**
- UNIQUE
 - default 224
 - unique index 287
 - creating 290
 - unique names
 - dataset names
 - avoiding 360
 - DDnames
 - avoiding 360
 - generating
 - multiple jobs for multiple activities 430
 - multiple jobs for single activity 430
 - unique setting
 - creating an index 289
 - unit value
 - for DASD estimators 118
 - UPDATE NOW option 233
 - user PDS 166
 - user variables
 - creating Help for 478
 - elements of template 470
 - Help template after modification 485
 - Help template before modification 484
 - prerequisites for creating 468
 - prerequisites for creating Help 479
 - references in template 472
 - template after modification 477
 - template before modification 476
 - template for creating 470
 - with utility profiles 464
 - user variables template panel
 - KTWPUTUS panel library member 469
 - user variables, creating 467

Index

- user variables, using 467
- Users
 - available authorizations 103
- using a different system PDS
 - authority 184
- using user variables 467
- UTIL command 355
 - as batch command 572
- Util (U) select 354
- UTILCMDS DD 539
- utilities
 - access from command line 355
 - access from utilities menu 355
 - accessing 344, 354
 - availability 344
 - batch 489
 - batch functions for 510
 - CHECK DATA utility 367
 - CHECK DATA utility with RI data 367
 - CHECK INDEX utility 369
 - Copy Data Sets (LOCAL) panel 438
 - Copy Data Sets (RECOVERY) panel 440
 - COPY utility 371
 - data for referential integrity 346
 - DB2
 - options 56
 - DIAG ABENDS utility 375
 - DIAG SYSUTIL utility 377
 - DISASTER REC utility 379
 - General Options panel 428
 - ICCAT utility 382
 - jcl for 521
 - List option 349
 - menus for 348
 - MERGE COPY utility 384
 - MOD parameter 361
 - MODIFY utility 387
 - on panels 344
 - online Help 423
 - options for DB2 commands 426
 - options for job control 428
 - values for options 428
 - processing statements 359
 - profile 350
 - QUIESCE utility 389
 - REBIND command 391
 - REBIND PKGS command 394
 - REBIND PLANS command 397
 - RECOVER INDEX utility 400
 - RECOVER utility for table spaces 402
- utilities (*continued*)
 - REORG INDEX utility 407
 - REORG utility for table spaces 409
 - REPORT utility 411
 - RUNSTATS utility 413
 - setting profiles 351
 - skeletons for 524
 - STOSPACE utility 417
 - using device (esoteric) names 445
 - using filtering with 354
 - Util (U) select 354
 - Utility Restart Options panel 431
 - utility stacking 361
 - values 356
 - values for options 426, 427
 - CMDS AFTER 427
 - CMDS BEFORE 426
 - Work Data Sets panel 432
 - utilities management
 - options 82
 - utilities menu
 - List option 349
 - profile option 452
 - utilities, batch
 - prerequisites 494
 - utilities, DB2
 - reviewing entries for accuracy 358
 - utility
 - profiles
 - creating 456
 - customization 456
 - modifying 456
 - utility commands
 - batch
 - AMS command 545
 - BIND command 547
 - DISP command 548
 - DUMP command 549
 - END command 552
 - EXEC/NOEXEC commands 554, 558
 - FLIST command 555
 - OUTC command 559
 - PDESCOPY command 562
 - PDSEEDIT command 564
 - PDSEXEC command 566
 - PDSLISIT command 567
 - PRNT command 569
 - START command 571
 - UTIL command 572
 - DUMP (batch)
 - DB operand 551

- utility generation
 - dataset sizing for 446
 - Utility ID
 - Display Utility Status
 - required values 82
 - Stop a Utility
 - required values 82
 - utility profile
 - accessing the menu 454
 - creating 456
 - customization 456
 - default 452
 - location for 452
 - modifying 456
 - options available 455
 - options available 453
 - possibility of multiple 452
 - providing values on the panels 460
 - saving 462
 - setting 453
 - specific 374
 - values for DIAG ABENDS utility 376
 - utility profile menu options
 - KTWPUOMU 458
 - utility profiles
 - with user variables 464
 - Utility Restart Options panel 431
 - utility skeleton, KTCUTJCL 452
 - utility, batch
 - introduction 530
 - printing reports 495
 - reports produced by 495
 - utility
 - profiles
 - providing values 460
- V**
- value migration
 - from the DSNZPARMs facility 127
 - values
 - CLOSE setting
 - creating a table space 257
 - completing for creating tables 281
 - creating an index
 - buffer pool name 289
 - CLOSE setting 289
 - clustering index online 289
 - Defer setting 289
 - subpage setting 289
 - values (*continued*)
 - creating an index (*continued*)
 - unique setting 289
 - creating storage groups
 - modifying from an existing storage group 312
 - using from an existing storage group 310
 - default
 - for creating indexes 296
 - defaults
 - creating a table space 260
 - filtering in batch 509
 - for column definition 278
 - for defaults 225
 - LOCKsize setting
 - creating a table space 257
 - object ID (OBID) 281
 - Partitions setting
 - creating a partitioned table space 257
 - using the default 260
 - providing for CREATE statements 220
 - SEGsize setting
 - creating a segmented table space 257
 - using the default 260
 - specifying
 - for job card field 429
 - for job statements field 429
 - utilities 356
 - VSAM dataset (index space)
 - ERASE setting 289
 - FREEPAGE value 289
 - Gbpcache value 289
 - PCTfree value 289
 - Priqty value 289
 - Secqty value 289
 - storage group name 289
 - VSAM password 289
 - values, available default 224
 - values, restoring default 223
 - VIEWNAME display field
 - FLIST command output panel 557
 - Views
 - available authorizations 103
 - VIO symbolic name 155
 - VOLSERS
 - for storage groups 308
 - VOLSERS list
 - editing 312

Index

VSAM Catalog 557
 default 224
VSAM datasets
 creating 216
 creating a table space
 prerequisites 252
 values 252
 DB2-created and managed 322
 defining and setting up 328
 defining at creation 215
 existing datasets
 prerequisites 325
 managing 292
 non-partitioned, DB2 managed 302
 partitions 320
 prerequisites for table spaces or index
 spaces 322
 setting up 326
 setting up DB2-created and managed 323
 setting up for table spaces 248
 storage group name 330
VSAM password
 creating a table space 252
 VSAM dataset values for index space 289

W

WAIT
 Archive Log
 required values 83
Work Data Sets panel 432
work dataset
 sizing for index recovery
 when object is missing 448
work datasets
 for sort processing
 allocation of 437

X

X (Exclude) select 298

Z

Z (Auths) 106
Z (Auths) select
 authority required to display 102
ZPARMs values
 authorization 128

User Comment Form

!DB®/WORKBENCH for DB2 Advanced User's Guide Version 500

TW54-5848-3

Please take a moment to share your comments and suggestions regarding Candle's documentation. Be as specific as possible.

Page #	Comment
①	_____
②	_____
③	_____
④	_____
⑤	_____
⑥	_____
⑦	_____

Please provide the following information. Thank you for your time and assistance.

Name

Job Function

Company

Site ID

Address

You may return your comments to us at your convenience. Upon receipt, we will respond to you within five business days.



FAX
(310) 727-4131




Fold, tape, and mail
(postage paid)
Please do not staple.

May we call to discuss your comments?

YES Phone number: _____ Best time to call: _____

NO


Cut or fold
along line




Fold and tape

Please do not staple

Fold and tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



BUSINESS REPLY MAIL
FIRST CLASS MAIL PERMIT NO 559 SANTA MONICA CA

POSTAGE WILL BE PAID BY ADDRESSEE



ATTN INFORMATION DEVELOPMENT
CANDLE CORPORATION
201 N DOUGLAS STREET
EL SEGUNDO CA 90245




Fold and tape

Please do not staple

Fold and tape

Cut or fold
along line





* T W 5 4 - 5 8 4 8 - 3 *