DB2 Universal Database for OS/390 and z/OS

Utility Guide and Reference

Version 7

SC26-9945-01

DB2 Universal Database for OS/390 and z/OS

Utility Guide and Reference

Version 7

SC26-9945-01

Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 725.

Second Edition, Softcopy Only (August 2001)

This edition applies to Version 7 of IBM DATABASE 2 Universal Database Server for OS/390 and z/OS (DB2 for OS/390 and z/OS), 5675-DB2, and to any subsequent releases until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

This softcopy version is based on the printed edition of the book and includes the changes indicated in the printed version by vertical bars. Additional changes made to this softcopy version of the book since the hardcopy book was published are indicated by the hash (#) symbol in the left-hand margin. Editorial changes that have no technical significance are not noted.

This and other books in the DB2 for OS/390 and z/OS library are periodically updated with technical changes. These updates are made available to licensees of the product on CD-ROM and on the Web (currently at www.ibm.com/software/data/db2/os390/library.html). Check these resources to ensure that you are using the most current information.

© Copyright International Business Machines Corporation 1983, 2001. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this book . . . xv
Who should read this book . . . xv
Conventions and terminology used in this book . . . xv
Product terminology and citations . . . xv
Naming conventions used in this book . . . xvi
How to read the syntax diagrams. . . . xviii
How to send your comments . . . xx

Summary of changes to this book . . . xxi

Part 1. Introduction . . . 1

Chapter 1. Summary of changes to DB2 for OS/390 and z/OS Version 7 . . . 3
Enhancements for managing data . . . 3
Enhancements for reliability, scalability, and availability. . . . 3
Easier development and integration of e-business applications . . . 4
Improved connectivity . . . 5
Features of DB2 for OS/390 and z/OS. . . . 6
Migration considerations . . . 6

Chapter 2. Introduction to the DB2 utilities . . . 7
Types of DB2 utilities . . . 7
Description of online utilities . . . 7
Description of stand-alone utilities . . . 7
Privileges and authorization IDs . . . 7
Running utilities on declared temporary tables . . . 8
Running utilities on objects with the DEFINE NO attribute . . . 8

Part 2. DB2 online utilities . . . 11

Chapter 3. Invoking DB2 online utilities . . . 23
Creating utility control statements . . . 23
Control statement coding rules . . . 23
Example of option description . . . 24
Data sets used by online utilities . . . 24
Concatenating data sets . . . 27
Controlling data set disposition . . . 27
Security . . . 28
Using the DB2 utilities panel in DB2I . . . 28
Using the DSNU CLIST command in TSO . . . 31
DSNU CLIST command syntax . . . 32
DSNU CLIST option descriptions . . . 32
Reviewing DSNU CLIST command output . . . 36
Editing the generated JCL data set . . . 37
Examples . . . 37
Using the supplied JCL procedure (DSNUPROC) . . . 38
DSNUPROC syntax . . . 38
DSNUPROC option descriptions . . . 38
Sample DSNUPROC listing . . . 40
Creating the JCL data set yourself. . . . 41
EXEC statement . . . 41

© Copyright IBM Corp. 1983, 2001

Chapter 4. Monitoring and controlling online utilities 43
Monitoring utilities with the DISPLAY UTILITY command 43
Determining the status of a utility 43
Determining which utility phase is currently executing. 44
Determining why a utility failed to complete 44
Running utilities concurrently. 44
Running online utilities in a data sharing environment. 45
Terminating an online utility with the TERM UTILITY command 45
Restarting an online utility. 46
Updating the JCL data set for restarting a utility. 46
Adding or deleting utility statements 47
Restarting after the output data set is full 47
RESTARTING with TEMPLATES 47
RESTARTING with LISTS 48
Other restart hints. 48

Chapter 5. CATMAINT 49
Syntax and options of the control statement 49
Syntax diagram. 49
Option descriptions 49
Instructions for running CATMAINT 49
Before running CATMAINT 50
Data sets used by CATMAINT 50
Instructions for specific tasks. 50
Terminating or restarting CATMAINT 50
Concurrency and compatibility 51

Chapter 6. CHECK DATA 53
Syntax and options of the control statement 54
Syntax diagram. 54
Option descriptions 54
Instructions for running CHECK DATA 58
Before running CHECK DATA 58
Data sets used by CHECK DATA 61
Creating the control statement 62
Instructions for specific tasks. 62
Terminating or restarting CHECK DATA 64
Concurrency and compatibility 65
Sample control statements 66

Chapter 7. CHECK INDEX 67
Syntax and options of the control statement 67
Syntax diagram. 68
Option descriptions 68
Instructions for running CHECK INDEX 69
Data sets used by CHECK INDEX. 70
Creating the control statement 71
Instructions for specific tasks. 71
Reviewing CHECK INDEX output 71
Terminating or restarting CHECK INDEX 72
Concurrency and compatibility 72
Sample control statements 73

Chapter 8. CHECK LOB 75
Syntax and options of the control statement 75
Syntax diagram. 76

Option descriptions . . . 76
Instructions for running CHECK LOB . . . 77
Before running CHECK LOB . . . 78
Data sets used by CHECK LOB . . . 78
Creating the control statement . . . 78
Instructions for specific tasks . . . 78
Terminating or restarting CHECK LOB . . . 79
Concurrency and compatibility . . . 79
Sample control statements . . . 80

Chapter 9. COPY . . . 81
Syntax and options of the control statement . . . 82
Syntax diagram . . . 82
Option descriptions . . . 84
Instructions for running COPY . . . 89
Before running COPY . . . 90
Data sets used by COPY . . . 90
Creating the control statement . . . 92
Instructions for specific tasks . . . 92
Considerations for running COPY . . . 100
Terminating or restarting COPY . . . 101
Concurrency and compatibility . . . 102
Sample control statements . . . 104

Chapter 10. COPYTOCOPY . . . 109
Syntax and options of the control statement . . . 110
Syntax diagram . . . 110
Option descriptions . . . 112
Instructions for running COPYTOCOPY . . . 115
Before running COPYTOCOPY . . . 115
Data sets used by COPYTOCOPY . . . 115
Creating the control statement . . . 116
Instructions for specific tasks . . . 116
Considerations for running COPYTOCOPY . . . 118
Terminating or restarting COPYTOCOPY . . . 119
Concurrency and compatibility . . . 119
Sample control statement . . . 120

Chapter 11. DIAGNOSE . . . 123
Syntax and options of the control statement . . . 123
Syntax diagram . . . 123
Option descriptions . . . 125
Instructions for running DIAGNOSE . . . 127
Data sets used by DIAGNOSE . . . 127
Instructions for specific tasks . . . 128
Terminating or restarting DIAGNOSE . . . 128
Concurrency and compatibility . . . 128
Sample control statements . . . 128

Chapter 12. EXEC SQL . . . 131
Syntax and options of the control statement . . . 131
Syntax diagram . . . 131
Concurrency and compatibility . . . 132
Sample control statements . . . 132

Chapter 13. LISTDEF . . . 135

Syntax and options of the control statement . . . 136
Syntax diagram . . . 136
Option descriptions . . . 138
Instructions for using LISTDEF . . . 143
Generating object lists using LISTDEF . . . 144
Concurrency and compatibility . . . 147
Sample control statements . . . 147

Chapter 14. LOAD . . . 151
Syntax and options of the control statement . . . 152
Syntax diagram . . . 153
Option descriptions . . . 155
Into-table-spec . . . 167
Option descriptions for INTO TABLE . . . 169
Instructions for running LOAD . . . 182
Before running LOAD . . . 183
Data sets used by LOAD . . . 183
Instructions for specific tasks . . . 186
Considerations for running LOAD . . . 198
Terminating or restarting LOAD . . . 204
Concurrency and compatibility . . . 206
After running LOAD . . . 208
Copying the loaded table space or partition . . . 209
Resetting the COPY-pending status . . . 209
Resetting the REBUILD-pending status . . . 209
Resetting the CHECK-pending status . . . 210
Collecting inline statistics while loading a table . . . 212
Running CHECK INDEX after loading a table having indexes . . . 212
Recovering a failed LOAD job . . . 212
Reorganizing an auxiliary index after LOAD . . . 212
Sample control statements . . . 213

Chapter 15. MERGECOPY . . . 221
Syntax and options of the control statement . . . 222
Syntax diagram . . . 222
Option descriptions . . . 222
Instructions for running MERGECOPY . . . 224
Data sets used by MERGECOPY . . . 225
Creating the control statement . . . 226
Instructions for specific tasks . . . 226
Terminating or restarting MERGECOPY . . . 228
Concurrency and compatibility . . . 228
Sample control statements . . . 229

Chapter 16. MODIFY RECOVERY . . . 231
Syntax and options of the control statement . . . 232
Syntax diagram . . . 232
Option descriptions . . . 232
Instructions for running MODIFY RECOVERY . . . 233
Before running MODIFY RECOVERY . . . 234
Data sets used by MODIFY RECOVERY . . . 234
Creating the control statement . . . 234
Instructions for specific tasks . . . 234
Terminating or restarting MODIFY RECOVERY . . . 235
Concurrency and compatibility . . . 235
Sample control statements . . . 236

Chapter 17. MODIFY STATISTICS 237
Syntax and options of the control statement 237
Syntax diagram 238
Option descriptions 238
Instructions for running MODIFY STATISTICS 239
Data sets used by MODIFY STATISTICS 240
Creating the control statement 240
Instructions for specific tasks 240
Terminating or restarting MODIFY STATISTICS 241
Concurrency and compatibility 241
Sample control statements 241

Chapter 18. OPTIONS 245
Syntax and options of the control statement 245
Syntax diagram 245
Option descriptions 246
Instructions for using OPTIONS 248
Concurrency and compatibility 249
Sample control statements 249

Chapter 19. QUIESCE 251
Syntax and options of the control statement 251
Syntax diagram 251
Option descriptions 252
Instructions for running QUIESCE 253
Before running QUIESCE 253
Data sets used by QUIESCE 253
Creating the control statement 254
Instructions for specific tasks 254
Considerations for running QUIESCE 255
Terminating or restarting QUIESCE 256
Concurrency and compatibility 256
Sample control statements 258

Chapter 20. REBUILD INDEX 259
Syntax and options of the control statement 259
Syntax diagram 259
Option descriptions 260
Instructions for running REBUILD INDEX 264
Before running REBUILD INDEX 264
Data sets used by REBUILD INDEX 264
Creating the control statement 265
Instructions for specific tasks 265
Terminating or restarting REBUILD INDEX 270
Concurrency and compatibility 271
Sample control statements 272

Chapter 21. RECOVER 275
Syntax and options of the control statement 276
Syntax diagram 276
Option descriptions 277
Instructions for running RECOVER 282
Before running RECOVER 282
Data sets used by RECOVER 283
Instructions for specific tasks 283
Considerations for running RECOVER 297

Allocating incremental image copies . . . 297
Performing fallback recovery . . . 298
Retaining tape mounts . . . 298
Avoiding damaged media . . . 299
Recovering table spaces and index spaces with mixed volume IDs . . . 299
Terminating or restarting RECOVER . . . 299
Concurrency and compatibility . . . 300
Sample control statements . . . 301

Chapter 22. REORG INDEX . . . 303
Syntax and options of the control statement . . . 304
Syntax diagram . . . 304
Option descriptions . . . 306
Instructions for running REORG INDEX . . . 314
Before running REORG INDEX . . . 314
Data sets used by REORG INDEX . . . 316
Creating the control statement . . . 316
Instructions for specific tasks . . . 316
Terminating or restarting REORG INDEX . . . 321
Concurrency and compatibility . . . 322
REORG INDEX compatibility . . . 323
Reviewing REORG INDEX output . . . 324
Sample control statements . . . 325

Chapter 23. REORG TABLESPACE . . . 327
Syntax and options of the control statement . . . 329
Syntax diagram . . . 329
Option descriptions . . . 335
REORG TABLESPACE options syntax . . . 355
Option descriptions for REORG TABLESPACE options . . . 355
Instructions for running REORG TABLESPACE . . . 356
Before running REORG TABLESPACE . . . 357
Data sets used by REORG TABLESPACE . . . 361
Creating the control statement . . . 364
Instructions for specific tasks . . . 365
Considerations for running REORG . . . 378
Terminating or restarting REORG TABLESPACE . . . 380
Concurrency and compatibility . . . 383
REORG TABLESPACE compatibility . . . 383
Reviewing REORG TABLESPACE output . . . 387
After running REORG TABLESPACE . . . 387
Sample control statements . . . 388

Chapter 24. REPAIR . . . 397
Syntax and options of the control statement . . . 398
Syntax diagram . . . 398
REPAIR option descriptions . . . 398
SET TABLESPACE and SET INDEX statement syntax . . . 399
SET TABLESPACE and SET INDEX option descriptions . . . 400
LOCATE block syntax . . . 401
LOCATE TABLESPACE statement option descriptions . . . 401
LOCATE INDEX statement option descriptions . . . 403
VERIFY statement syntax . . . 403
VERIFY statement option descriptions . . . 404
REPLACE statement syntax . . . 404
REPLACE statement option descriptions . . . 405

DELETE statement syntax and description . . . 405
DUMP statement syntax . . . 406
DUMP statement option descriptions . . . 406
DBD statement syntax . . . 407
DBD statement option descriptions . . . 408
Instructions for running REPAIR . . . 409
Before running REPAIR . . . 409
Data sets used by REPAIR . . . 410
Creating the control statement . . . 410
Instructions for specific tasks . . . 410
Terminating or restarting REPAIR . . . 413
Concurrency and compatibility . . . 413
Reviewing REPAIR output . . . 415
After running REPAIR . . . 416
Sample control statements . . . 416

Chapter 25. REPORT . . . 419
Syntax and options of the control statement . . . 420
Syntax diagram . . . 420
Option descriptions . . . 420
Instructions for running REPORT . . . 423
Data sets used by REPORT . . . 423
Creating the control statement . . . 424
Instructions for specific tasks . . . 424
Terminating or restarting REPORT . . . 425
Concurrency and compatibility . . . 425
Reviewing REPORT output . . . 426
Sample control statements . . . 429

Chapter 26. RUNSTATS . . . 431
Syntax and options of the control statement . . . 432
RUNSTATS TABLESPACE syntax diagram . . . 432
RUNSTATS TABLESPACE option descriptions . . . 434
RUNSTATS INDEX syntax diagram . . . 438
RUNSTATS INDEX option descriptions . . . 438
Instructions for running RUNSTATS . . . 441
Before running RUNSTATS . . . 441
Data sets used by RUNSTATS . . . 442
Creating the control statement . . . 442
Instructions for specific tasks . . . 442
Terminating or restarting RUNSTATS . . . 443
Concurrency and compatibility . . . 444
Reviewing RUNSTATS output . . . 445
Access path statistics . . . 447
Space statistics (columns for tuning information) . . . 449
After running RUNSTATS . . . 453
Sample control statements . . . 453

Chapter 27. STOSPACE . . . 457
Syntax and options of the control statement . . . 457
Syntax diagram . . . 457
Option descriptions . . . 457
Instructions for running STOSPACE . . . 458
Data sets used by STOSPACE . . . 458
Creating the control statement . . . 458
Instructions for specific tasks . . . 459

Considerations for running STOSPACE 460
Terminating or restarting STOSPACE 460
Concurrency and compatibility 460
Reviewing STOSPACE output 461
Sample control statement 461

Chapter 28. TEMPLATE 463
Syntax and options of the control statement 463
Syntax diagram 463
Option descriptions 464
Instructions for using TEMPLATE 472
Key TEMPLATE operations 473
Default space calculations 474
Working with TAPE 474
Working with GDGs 475
Concurrency and compatibility 475
Sample control statements 475

Chapter 29. UNLOAD 477
Syntax and options of the control statement 477
Syntax diagram 477
Option descriptions 478
FROM-TABLE-spec 485
Option descriptions for FROM TABLE 489
Instructions for running UNLOAD 507
Data sets used by UNLOAD 508
Instructions for specific tasks 508
Terminating or restarting UNLOAD 518
Concurrency and compatibility 518
Sample control statements 520

Part 3. Stand-alone utilities 523

Chapter 30. Invoking stand-alone utilities 527
Creating utility statements and EXEC PARM parameters 527
Creating utility control statements 527
Specifying options using the JCL EXEC PARM parameter 527
Example of an option description 528

Chapter 31. DSNJLOGF (Preformat Active Log) 529
Before running DSNJLOGF 529
Environment 529
Control statement 529
Sample control statement 529
DSNJLOGF output 529

Chapter 32. DSNJU003 (Change Log Inventory) 531
Syntax and options of the control statement 531
DSNJU003 (change log inventory) syntax diagram 531
Option descriptions 533
Before running DSNJU003 539
Environment 539
Authorization required 540
Control statement 540
Using DSNJU003 541
Running DSNJU003 541

Making changes for active logs . . . 541
Making changes for archive logs . . . 543
Creating a conditional restart control record . . . 543
Deleting log data sets with errors . . . 543
Altering references to NEWLOG and DELETE data sets . . . 544
Specifying the NEWCAT statement . . . 544
Renaming DB2 system data sets . . . 545
Renaming DB2 active log data sets . . . 545
Renaming DB2 archive log data sets . . . 546
Sample control statements . . . 546

Chapter 33. DSNJU004 (Print Log Map) . . . 547
Syntax and options of the control statement . . . 547
DSNJU004 (print log map) syntax diagram . . . 547
Option descriptions . . . 547
Before running DSNJU004 . . . 548
Environment . . . 548
Authorization required . . . 548
Control statement . . . 548
Recommendations . . . 549
Sample control statement . . . 549
DSNJU004 (Print Log Map) output . . . 549
Timestamps in the BSDS . . . 550
Active log data set status . . . 551
Reading conditional restart control records . . . 554

Chapter 34. DSN1CHKR . . . 555
Syntax and options of the control statement . . . 555
DSN1CHKR syntax diagram . . . 555
Option descriptions . . . 555
Before running DSN1CHKR . . . 557
Environment . . . 557
Authorization required . . . 557
Control statement . . . 557
Restrictions . . . 557
Sample control statements . . . 558
DSN1CHKR output . . . 561

Chapter 35. DSN1COMP . . . 563
Syntax and options of the control statement . . . 563
DSN1COMP syntax diagram . . . 563
Option descriptions . . . 563
Before running DSN1COMP . . . 566
Environment . . . 566
Authorization required . . . 566
Control statement . . . 566
Recommendation . . . 567
Using DSN1COMP . . . 567
Estimating compression savings achieved by REORG . . . 567
Including free space in compression calculations . . . 568
Running DSN1COMP on a table space with identical data . . . 568
Sample control statements . . . 568
DSN1COMP output . . . 568
Message DSN1941 . . . 568
Sample DSN1COMP report . . . 569

Chapter 36. DSN1COPY 571
Syntax and options of the control statement 571
DSN1COPY syntax diagram 572
Option descriptions 572
Before running DSN1COPY 576
Environment 576
Authorization required 576
Control statement 577
Defining the input data set 579
Defining the output data set 580
Restrictions 581
Recommendations 581
Using DSN1COPY 582
Altering a table before running DSN1COPY 582
Checking for inconsistent data 582
Translating DB2 internal identifiers 582
Using an image copy as input to DSN1COPY 583
Resetting page log RBAs 583
Copying multiple data set table spaces 583
Restoring indexes with DSN1COPY 583
Restoring table spaces with DSN1COPY 584
Printing with DSN1COPY 584
Copying tables from one subsystem to another 584
Sample control statements 585
DSN1COPY output 586

Chapter 37. DSN1LOGP 587
Syntax and options of the control statement 587
DSN1LOGP syntax diagram 588
Option descriptions 588
Before running DSN1LOGP 593
Environment 593
Authorization required 594
Control statement 594
Using DSN1LOGP 595
Reading archive log data sets on tape 595
Locating table and index identifiers 596
Sample control statements 596
DSN1LOGP output 599
Reviewing DSN1LOGP output 599
Interpreting error codes 607

Chapter 38. DSN1PRNT 609
Syntax and options of the control statement 609
DSN1PRNT syntax diagram 610
Option descriptions 610
Before running DSN1PRNT 614
Environment 614
Authorization required 614
Control statement 614
Recommendations 615
Sample control statements 615
DSN1PRNT output 616

Chapter 39. DSN1SDMP 617
Syntax and options of the control statement 617

DSN1SDMP syntax diagram . . . 617
Option descriptions . . . 617
Before running DSN1SDMP . . . 620
Environment . . . 620
Authorization required . . . 620
Control statement . . . 621
Using DSN1SDMP . . . 621
Assigning buffers . . . 621
Generating a dump . . . 622
Stopping or modifying DSN1SDMP traces . . . 622
Sample control statements . . . 623
DSN1SDMP output . . . 624

Part 4. Appendixes . . . 625

Appendix A. Limits in DB2 for OS/390 and z/OS . . . 627

Appendix B. Stored procedures shipped with DB2 . . . 631
Invoking utilities as a stored procedure (DSNUTILS) . . . 631
Environment . . . 632
Authorization required . . . 632
Control statement . . . 632
DSNUTILS syntax diagram . . . 634
DSNUTILS option descriptions . . . 634
Modifying the WLM-established address space . . . 641
Sample program for calling DSNUTILS . . . 641
DSNUTILS output . . . 641
The DB2 UDB Control Center table space and index information stored procedure (DSNACCQC) . . . 642
Environment . . . 642
Authorization required . . . 642
DSNACCQC syntax diagram . . . 642
DSNACCQC option descriptions . . . 643
Example of DSNACCQC invocation . . . 647
DSNACCQC output . . . 647
The DB2 UDB Control Center partition information stored procedure (DSNACCAV) . . . 649
Environment . . . 649
Authorization required . . . 649
DSNACCAV syntax diagram . . . 650
DSNACCAV option descriptions . . . 650
Example of DSNACCAV invocation . . . 652
DSNACCAV output . . . 653
The DB2 real-time statistics stored procedure . . . 658
Environment . . . 659
Authorization required . . . 659
DSNACCOR syntax diagram . . . 660
DSNACCOR option descriptions . . . 660
Formulas for recommending actions . . . 667
Using an exception table . . . 668
Example of DSNACCOR invocation . . . 669
DSNACCOR output . . . 673

Appendix C. Resetting an advisory or restrictive status . . . 677
Auxiliary CHECK-pending status . . . 677
Auxiliary warning status . . . 677

CHECK-pending status 678
COPY-pending status 679
Group buffer pool RECOVER-pending status 679
Informational COPY-pending status 680
REBUILD-pending status 680
RECOVER-pending status 681
REFRESH-pending status 682
REORG-pending status 682
Restart pending status 683

Appendix D. How to run sample programs DSNTIAUL, DSNTIAD, and DSNTEP2 685
Running DSNTIAUL 686
Running DSNTIAD 689
Running DSNTEP2 691

Appendix E. DB2 utilities packaging 695
SMP/E jobs for DB2 utility products 696
Installing DB2 Operational Utilities 696
Installing DB2 Diagnostic & Recovery Utilities 696
Installing DB2 Utilities Suite 696
How to operate in a mixed-release data sharing environment 697

Appendix F. Real-time statistics tables 699
Setting up your system for real-time statistics 699
Creating and altering the real-time statistics objects 699
Setting the interval for writing real-time statistics 700
Starting the real-time statistics database 701
Contents of the real-time statistics tables 701
Operating with real-time statistics 713
When DB2 externalizes real-time statistics 713
How DB2 utilities affect the real-time statistics 714
How non-DB2 utilities affect real-time statistics 720
Real-time statistics on objects in work file databases and the TEMP database 721
Real-time statistics on read-only objects 721
How dropping objects affects real-time statistics 721
How SQL operations affect real-time statistics counters 721
Real-time statistics in data sharing 722
Improving concurrency with real-time statistics 722
Recovering the real-time statistics tables 722
Statistics accuracy 722

Notices 725
Programming interface information 726
Trademarks 727

Glossary 729

Bibliography 745

Index 751

About this book

This book contains usage information for the tasks of system administration, database administration, and operation. It presents detailed information on using utilities, specifying syntax (including keyword and parameter descriptions), and starting, stopping, and restarting utilities. This book also includes Job control language (JCL) and control statements for each utility.

Important

In this version of DB2® for OS/390® and z/OS, some utility functions are available as optional products. You must separately order and purchase a license to such utilities, and discussion of those utility functions in this publication is not intended to otherwise imply that you have a license to them.

Who should read this book

This book is intended for system administrators, database administrators, system operators, and application programmers of DB2 online and stand-alone utilities. We recommend that you familiarize yourself with DB2 for OS/390 and z/OS prior to using this book.

Conventions and terminology used in this book

This section provides information about conventions in this book.

Product terminology and citations

In this book, DB2 Universal Database™ Server for OS/390 and z/OS is referred to as "DB2 for OS/390 and z/OS." In cases where the context makes the meaning clear, DB2 for OS/390 and z/OS is referred to as "DB2." When this book refers to other books in this library, a short title is used. (For example, "See *DB2 SQL Reference*" is a citation to *IBM® DATABASE 2™ Universal Database Server for OS/390 and z/OS SQL Reference*.)

When referring to a DB2 product other than DB2 for OS/390 and z/OS, this book uses the product's full name to avoid ambiguity.

The following terms are used as indicated:

DB2 Represents either the DB2 licensed program or a particular DB2 subsystem.

C and C language
Represent the C programming language.

CICS® Represents CICS/ESA® and CICS Transaction Server for OS/390.

IMS™ Represents IMS or IMS/ESA®.

MVS Represents the MVS element of OS/390.

OS/390
Represents the OS/390 or z/OS operating system.

© Copyright IBM Corp. 1983, 2001

RACF®

Represents the functions that are provided by the RACF component of the SecureWay® Security Server for OS/390 or by the RACF component of the OS/390 Security Server.

Naming conventions used in this book

This section describes naming conventions unique to commands and utilities.

When a parameter refers to an object created by SQL statements (for example, tables, table spaces, and indexes), follow the SQL syntactical naming conventions.

Characters are classified as *letters, digits,* or *special characters*.

- A *letter* is any one of the uppercase characters A through Z (including the three characters reserved as alphabetic extenders for national languages, #, @, and $ in the United States).
- A *digit* is any one of the characters 0 through 9.
- A *special character* is any character other than a letter or a digit.

See Chapter 2 of *DB2 SQL Reference* for an additional explanation of long identifiers, short identifiers, and location identifiers.

authorization-id

A short identifier of 1 to 8 letters, digits, or the underscore that identifies a set of privileges. An authorization ID must begin with a letter.

connection-name

An identifier of 1 to 8 characters that identifies an address space connection to DB2. A connection identifier is one of the following:

- For DSN processes running in TSO foreground, the connection name TSO is used.
- For DSN processes running in TSO batch, the connection name BATCH is used.
- For the call attachment facility (CAF), the connection name DB2CALL is used.
- For IMS and CICS processes, the connection name is the system identification name.

See Part 4 (Volume 1) of *DB2 Administration Guide* for more information about managing DB2 connections.

correlation-id

An identifier of 1 to 12 characters that identifies a process within an address space connection. A correlation ID must begin with a letter.

A correlation ID can be one of the following:

- For DSN processes running in TSO foreground, the correlation ID is the TSO logon identifier.
- For DSN processes running in TSO batch, the correlation ID is the job name.
- For CAF processes, the correlation ID is the TSO logon identifier.
- For IMS processes, the correlation ID is the PST#.PSBNAME.
- For CICS processes, the correlation ID is the entry identifier.thread_number.transaction_identifier.

See Part 4 (Volume 1) of *DB2 Administration Guide* for more information about correlation IDs.

database-name
A short identifier that designates a database. The identifier must start with a letter and must not include special characters.

data-set-name
An identifier of 1 to 44 characters that identifies a data set.

dbrm-member-name
An identifier of 1 to 8 letters or digits that identifies a member of a partitioned data set.

A DBRM member name should not begin with DSN; this can sometimes conflict with DB2-provided DBRM member names. If a DBRM member name beginning with DSN is specified, DB2 issues a warning message.

dbrm-pds-name
An identifier of 1 to 44 characters that identifies a partitioned data set.

ddname
An identifier of 1 to 8 characters that designates the name of a DD statement.

hexadecimal-constant
A sequence of digits or any of the letters from A to F (uppercase or lowercase).

hexadecimal-string
An X followed by a sequence of characters that begins and ends with an apostrophe. The characters between the string delimiters must be a hexadecimal number.

index-name
A qualified or unqualified name that designates an index.

A qualified index name is a short identifier followed by a period and a long identifier. The short identifier is the authorization ID that owns the index.

An unqualified index name is a long identifier with an implicit qualifier. The implicit qualifier is an authorization ID that is determined by the rules set forth in Chapter 2 of *DB2 SQL Reference*.

location-name
A location identifier of 1 to 16 letters (but excluding the alphabetic extenders), digits or the underscore that identifies an instance of a data base management system. A location name must begin with a letter.

luname
An SQL short identifier of 1 to 8 characters that identifies a logical unit name. An luname must begin with a letter.

member-name
or digits that identifies a member of a partitioned data set.

A member name should not begin with DSN; this can sometimes conflict with DB2-provided member names. If a member name beginning with DSN is specified, DB2 issues a warning message.

qualifier-name
An SQL short identifier of 1 to 8 letters, digits, or the underscore that identifies the implicit qualifier for unqualified table names, views, indexes, and aliases.

string
A sequence of characters that begins and ends with an apostrophe.

subsystem-name
An identifier that specifies the DB2 subsystem as it is known to MVS.

table-name
A qualified or unqualified name that designates a table.

A fully qualified table name is a three-part name. The first part is a location name that designates the DBMS at which the table is stored. The second part is the authorization ID that designates the owner of the table. The third part is a long identifier. A period must separate each of the parts.

A two-part table name is implicitly qualified by the location name of the current server. The first part is the authorization ID that designates the owner of the table. The second part is a SQL long identifier. A period must separate the two parts.

A one-part or unqualified table name is a long identifier with two implicit qualifiers. The first implicit qualifier is the location name of the current server. The second is an authorization ID, which is determined by the rules set forth in Chapter 2 of *DB2 SQL Reference*.

table-space-name
A short identifier that designates a table space of an identified database. The identifier must start with a letter and must not include special characters. If a database is not identified, a table space name specifies a table space of database DSNDB04.

utility-id
An identifier of 1 to 16 characters that uniquely identifies a utility process within DB2. A utility ID must begin with a letter. The remaining characters can be upper and lower case letters, numbers 0 through 9, and the following characters: #, $, @, ¢, !, ¬, and ..

How to read the syntax diagrams

The following rules apply to the syntax diagrams used in this book:

- Read the syntax diagrams from left to right, from top to bottom, following the path of the line.

  The ►►— symbol indicates the beginning of a statement.

  The —► symbol indicates that the statement syntax is continued on the next line.

  The ►— symbol indicates that a statement is continued from the previous line.

  The —►◄ symbol indicates the end of a statement.
- Required items appear on the horizontal line (the main path).

- Optional items appear below the main path.

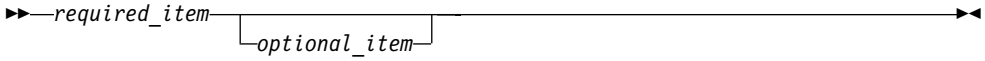

  If an optional item appears above the main path, that item has no effect on the execution of the statement and is used only for readability.

- If you can choose from two or more items, they appear vertically, in a stack.

  If you *must* choose one of the items, one item of the stack appears on the main path.

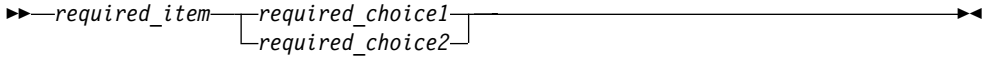

  If choosing one of the items is optional, the entire stack appears below the main path.

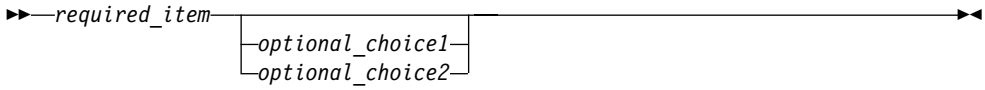

  If one of the items is the default, it appears above the main path and the remaining choices are shown below.

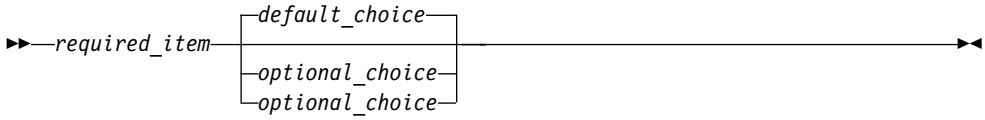

- An arrow returning to the left, above the main line, indicates an item that can be repeated.

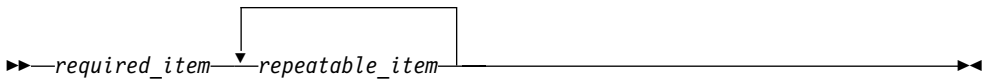

  If the repeat arrow contains a comma, you must separate repeated items with a comma.

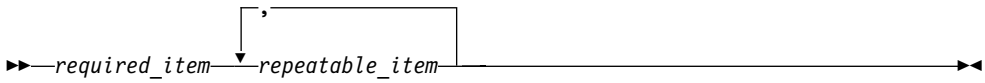

  A repeat arrow above a stack indicates that you can repeat the items in the stack.
- Keywords appear in uppercase (for example, `FROM`). They must be spelled exactly as shown. Variables appear in all lowercase letters (for example, *column-name*). They represent user-supplied names or values.
- If punctuation marks, parentheses, arithmetic operators, or other such symbols are shown, you must enter them as part of the syntax.

How to send your comments

Your feedback helps IBM to provide quality information. Please send any comments that you have about this book or other DB2 for OS/390 and z/OS documentation. You can use any of the following methods to provide comments:

- Send your comments by e-mail to db2pubs@vnet.ibm.com and include the name of the product, the version number of the product, and the number of the book. If you are commenting on specific text, please list the location of the text (for example, a chapter and section title, page number, or a help topic title).
- Send your comments from the Web. Visit the Web site at:

  http://www.ibm.com/software/db2os390

  The Web site has a feedback page that you can use to send comments.
- Complete the readers' comment form at the back of the book and return it by mail, by fax (800-426-7773 for the United States and Canada), or by giving it to an IBM representative.

Summary of changes to this book

The following online utilities and utility control statements are new for Version 7:

"Chapter 10. COPYTOCOPY" on page 109
"Chapter 12. EXEC SQL" on page 131
"Chapter 13. LISTDEF" on page 135
"Chapter 17. MODIFY STATISTICS" on page 237
"Chapter 18. OPTIONS" on page 245
"Chapter 28. TEMPLATE" on page 463
"Chapter 29. UNLOAD" on page 477

The following online utilities have changed:

"Chapter 6. CHECK DATA" on page 53
"Chapter 7. CHECK INDEX" on page 67
"Chapter 8. CHECK LOB" on page 75
"Chapter 9. COPY" on page 81
"Chapter 14. LOAD" on page 151
"Chapter 15. MERGECOPY" on page 221
"Chapter 16. MODIFY RECOVERY" on page 231
"Chapter 19. QUIESCE" on page 251
"Chapter 20. REBUILD INDEX" on page 259
"Chapter 21. RECOVER" on page 275
"Chapter 22. REORG INDEX" on page 303
"Chapter 23. REORG TABLESPACE" on page 327
"Chapter 25. REPORT" on page 419
"Chapter 26. RUNSTATS" on page 431

The following stand-alone utilities have minor changes:

"Chapter 34. DSN1CHKR" on page 555
"Chapter 35. DSN1COMP" on page 563
"Chapter 36. DSN1COPY" on page 571
"Chapter 38. DSN1PRNT" on page 609

The following appendix is new for Version 7:

- "Appendix E. DB2 utilities packaging" on page 695

© Copyright IBM Corp. 1983, 2001

Part 1. Introduction

Chapter 1. Summary of changes to DB2 for OS/390 and z/OS Version 7 . . . 3
Enhancements for managing data . . . 3
Enhancements for reliability, scalability, and availability . . . 3
Easier development and integration of e-business applications . . . 4
Improved connectivity . . . 5
Features of DB2 for OS/390 and z/OS . . . 6
Migration considerations . . . 6

Chapter 2. Introduction to the DB2 utilities . . . 7
Types of DB2 utilities . . . 7
Description of online utilities . . . 7
Description of stand-alone utilities . . . 7
Privileges and authorization IDs . . . 7
Running utilities on declared temporary tables . . . 8
Running utilities on objects with the DEFINE NO attribute . . . 8

© Copyright IBM Corp. 1983, 2001

Chapter 1. Summary of changes to DB2 for OS/390 and z/OS Version 7

DB2 for OS/390 and z/OS Version 7 delivers an enhanced relational database server solution for OS/390. This release focuses on greater ease and flexibility in managing your data, better reliability, scalability, and availability, and better integration with the DB2 family.

In Version 7, some utility functions are available as optional products; you must separately order and purchase a license to such utilities. Discussion of utility functions in this publication is not intended to otherwise imply that you have a license to them. See *DB2 Utility Guide and Reference* for more information about utilities products.

Enhancements for managing data

Version 7 delivers the following enhancements for managing data:

- DB2 now collects a comprehensive statistics history that:
  - Lets you track changes to the physical design of DB2 objects
  - Lets DB2 predict future space requirements for table spaces and indexes more accurately and run utilities to improve performance
- Database administrators can now manage DB2 objects more easily and no longer must maintain their utility jobs (even when new objects are added) by using enhancements that let them:
  - Dynamically create object lists from a pattern-matching expression
  - Dynamically allocate the data sets that are required to process those objects
- More flexible DBADM authority lets database administrators create views for other users.
- Enhancements to management of constraints let you specify a constraint at the time you create primary or unique keys. A new restriction on the DROP INDEX statement requires that you drop the primary key, unique key, or referential constraint before you drop the index that enforces a constraint.

Enhancements for reliability, scalability, and availability

Version 7 delivers the following enhancements for the reliability, scalability, and availability of your e-business:

- The DB2 Utilities Suite provides utilities for all of your data management tasks that are associated with the DB2 catalog.
- The new UNLOAD utility lets you unload data from a table space or an image copy data set. In most cases, the UNLOAD utility is faster than the DSNTIAUL sample program, especially when you activate partition parallelism for a large partitioned table space. UNLOAD is also easier to use than REORG UNLOAD EXTERNAL.
- The new COPYTOCOPY utility lets you make additional image copies from a primary image copy and registers those copies in the DB2 catalog. COPYTOCOPY leaves the target object in read/write access mode (UTRW), which allows Structured Query Language (SQL) statements and some utilities to run concurrently with the same target objects.

© Copyright IBM Corp. 1983, 2001

- Parallel LOAD with multiple inputs lets you easily load large amounts of data into partitioned table spaces for use in data warehouse applications or business intelligence applications. Parallel LOAD with multiple inputs runs in a single step, rather than in different jobs.
- A faster online REORG is achieved through the following enhancements:
  - Online REORG no longer renames data sets, which greatly reduces the time that data is unavailable during the SWITCH phase.
  - Additional parallel processing improves the elapsed time of the BUILD2 phase of REORG SHRLEVEL(CHANGE) or SHRLEVEL(REFERENCE).
- More concurrency with online LOAD RESUME is achieved by letting you give users read and write access to the data during LOAD processing so that you can load data concurrently with user transactions.
- More efficient processing for SQL queries:
  - More transformations of subqueries into a join for some UPDATE and DELETE statements
  - Fewer sort operations for queries that have an ORDER BY clause and WHERE clauses with predicates of the form COL=*constant*
  - More parallelism for IN-list index access, which can improve performance for queries involving IN-list index access
- The ability to change system parameters without stopping DB2 supports online transaction processing and e-business without interruption.
- Improved availability of user objects that are associated with failed or canceled transactions:
  - You can cancel a thread without performing rollback processing.
  - Some restrictions imposed by the restart function have been removed.
  - A NOBACKOUT option has been added to the CANCEL THREAD command.
- Improved availability of the DB2 subsystem when a log-read failure occurs: DB2 now provides a timely warning about failed log-read requests and the ability to retry the log read so that you can take corrective action and avoid a DB2 outage.
- Improved availability in the data sharing environment:
  - Group attachment enhancements let DB2 applications generically attach to a DB2 data sharing group.
  - A new LIGHT option of the START DB2 command lets you restart a DB2 data sharing member with a minimal storage footprint, and then terminate normally after DB2 frees the retained locks that it can.
  - You can let changes in structure size persist when you rebuild or reallocate a structure.
- Additional data sharing enhancements include:
  - Notification of incomplete units of recovery
  - Use of a new OS/390 and z/OS function to improve failure recovery of group buffer pools
- An additional enhancement for e-business provides improved performance with preformatting for INSERT operations.

Easier development and integration of e-business applications

Version 7 provides the following enhancements, which let you more easily develop and integrate applications that access data from various DB2 operating systems and distributed environments:

- DB2 XML Extender for OS/390 and z/OS, a new member of the DB2 Extender family, lets you store, retrieve, and search XML documents in a DB2 database.

- Improved support for UNION and UNION ALL operators in a view definition, a nested table expression, or a subquery predicate, improves DB2 family compatibility and is consistent with SQL99 standards.
- More flexibility with SQL gives you greater compatibility with DB2 on other operating systems:
  - Scrollable cursors let you move forward, backward, or randomly through a result table or a result set. You can use scrollable cursors in any DB2 applications that do not use DB2 private protocol access.
  - A search condition in the WHERE clause can include a subquery in which the base object of both the subquery and the searched UPDATE or DELETE statement are the same.
  - A new SQL clause, FETCH FIRST *n* ROWS, improves performance of applications in a distributed environment.
  - Fast implicit close in which the DB2 server, during a distributed query, automatically closes the cursor when the application attempts to fetch beyond the last row.
  - Support for options USER and USING in a new authorization clause for CONNECT statements lets you easily port applications that are developed on the workstation to DB2 for OS/390 and z/OS. These options also let applications that run under WebSphere to reuse DB2 connections for different users and to enable DB2 for OS/390 and z/OS to check passwords.
  - For positioned updates, you can specify the FOR UPDATE clause of the cursor SELECT statement without a list of columns. As a result, all updatable columns of the table or view that is identified in the first FROM clause of the fullselect are included.
  - A new option of the SELECT statement, ORDER BY *expression*, lets you specify operators as the sort key for the result table of the SELECT statement.
  - New datetime ISO functions return the day of the week with Monday as day 1 and every week with seven days.
- Enhancements to Open Database Connectivity (ODBC) provide partial ODBC 3.0 support, including many new application programming interfaces (APIs), which increase application portability and alignment with industry standards.
- Enhancements to the LOAD utility let you load the output of an SQL SELECT statement directly into a table.
- A new component called Precompiler Services lets compiler writers modify their compilers to invoke Precompiler Services and produce an *SQL statement coprocessor*. An SQL statement coprocessor performs the same functions as the DB2 precompiler, but it performs those functions at compile time. If your compiler has an SQL statement coprocessor, you can eliminate the precompile step in your batch program preparation jobs for COBOL and PL/I programs.
- Support for Unicode-encoded data lets you easily store multilingual data within the same table or on the same DB2 subsystem. The Unicode encoding scheme represents the code points of many different geographies and languages.

Improved connectivity

Version 7 offers improved connectivity:

- Support for COMMIT and ROLLBACK in stored procedures lets you commit or roll back an entire unit of work, including uncommitted changes that are made from the calling application before the stored procedure call is made.

- Support for Windows Kerberos security lets you more easily manage workstation clients who seek access to data and services from heterogeneous environments.
- Global transaction support for distributed applications lets independent DB2 agents participate in a global transaction that is coordinated by an XA-compliant transaction manager on a workstation or a gateway server (Microsoft Transaction Server or Encina, for example).
- Support for a DB2 Connect Version 7 enhancement lets remote workstation clients quickly determine the amount of time that DB2 takes to process a request (the server elapsed time).
- Additional enhancements include:
  - Support for connection pooling and transaction pooling for IBM DB2 Connect
  - Support for DB2 Call Level Interface (DB2 CLI) bookmarks on DB2 UDB for UNIX, Windows, OS/2

Features of DB2 for OS/390 and z/OS

Version 7 of DB2 UDB Server for OS/390 and z/OS offers several features that help you integrate, analyze, summarize, and share data across your enterprise:

- DB2 Warehouse Manager feature. The DB2 Warehouse Manager feature brings together the tools to build, manage, govern, and access DB2 for OS/390 and z/OS-based data warehouses. The DB2 Warehouse Manager feature uses proven technologies with new enhancements that are not available in previous releases, including:
  - DB2 Warehouse Center, which includes:
    - DB2 Universal Database Version 7 Release 1 Enterprise Edition
    - Warehouse agents for UNIX, Windows, and OS/390
    - Information Catalog
  - QMF Version 7
  - QMF High Performance Option
  - QMF for Windows
- DB2 Management Clients Package. The elements of the DB2 Management Clients Package are:
  - DB2 Control Center
  - DB2 Stored Procedure Builder
  - DB2 Installer
  - DB2 Visual Explain
  - DB2 Estimator
- Net Search Extender for in-memory text search for e-business applications
- Net.Data for secure Web applications

Migration considerations

Migration with full fallback protection is available when you have either DB2 for OS/390 Version 5 or Version 6 installed. You should ensure that you are fully operational on DB2 for OS/390 Version 5, or later, before migrating to DB2 for OS/390 and z/OS Version 7.

To learn about all of the migration considerations from Version 5 to Version 7, read the *DB2 Release Planning Guide* for Version 6 and Version 7; to learn about content information, also read appendixes A through F in both books.

Chapter 2. Introduction to the DB2 utilities

This chapter provides an introduction to the DB2 online and stand-alone utilities. This chapter also discusses the authorization rules for coding utility control statements and the data sets used by utilities.

Types of DB2 utilities

There are two types of DB2 utilities: online utilities and stand-alone utilities. Both types are described in this section.

Description of online utilities

DB2 online utilities run as standard MVS batch jobs or stored procedures, and they require DB2 to be running. They do not run under control of the terminal monitor program (TMP), but have their own attach mechanisms. They invoke DB2 control facility services directly. Refer to “Chapter 3. Invoking DB2 online utilities” on page 23 for information about the ways to execute these utilities.

Description of stand-alone utilities

The stand-alone utilities execute as batch jobs independent of DB2. They can be executed only by means of MVS JCL. Refer to the chapters on the individual utilities in “Chapter 30. Invoking stand-alone utilities” on page 527 for information about the ways to execute these utilities.

Privileges and authorization IDs

The issuer of a command or a utility job can be an individual user. It can also be a program running in batch mode or an IMS or CICS transaction. We use the term *process* to represent any or all of those.

A process is represented to DB2 by a set of identifiers (IDs). What the process can do with DB2 is determined by *privileges* and *authorities* that can be held by its identifiers. We use *“the privilege set of a process”* to mean the entire set of privileges and authorities that can be used by the process in a specific situation.

Three types of identifiers exist: *primary authorization IDs*, *secondary authorization IDs,* and *SQL authorization IDs.*

- Generally, the primary authorization ID identifies a specific process. For example, in the process initiated through the TSO attachment facility, the primary authorization ID is identical to the TSO logon ID. A trace record identifies the process by that ID.
- Secondary authorization IDs, which are optional, can hold additional privileges available to the process. A secondary authorization ID is often a Resource Access Control Facility (RACF) group ID. For example, a process can belong to a RACF group that holds the LOAD privilege on a particular database. Any member of the group can run the LOAD utility to load table spaces in the database.

  DB2 commands entered from an MVS console are not associated with any secondary authorization IDs.
- An SQL authorization ID (SQL ID) holds the privileges exercised when issuing certain dynamic SQL statements. This ID plays little part in the utilities described in this book.

© Copyright IBM Corp. 1983, 2001

Within DB2, a process can be represented by a primary authorization ID and possibly one or more secondary IDs. For detailed instructions on how to associate a process with one or more IDs, and how to grant privileges to those IDs, see "Processing connections" and "Processing sign-ons" in Part 3 (Volume 1) of *DB2 Administration Guide*.

A privilege or authority is granted to, or revoked from, an identifier by executing an SQL GRANT or a REVOKE statement. For the complete syntax of those statements, see Chapter 5 of *DB2 SQL Reference*.

If you use the access control authorization exit, then that exit might control the authorization rules, rather than the exits documented for each utility.

Running utilities on declared temporary tables

DB2 utilities operate on a variety of target objects. You can use the REPAIR DBD utility on declared temporary tables, which must be created in a database that is defined with the AS TEMP clause. No other DB2 utilities can be used on a declared temporary table, its indexes, or its table spaces.

For detailed information about target object support, see the "Concurrency and compatibility" section in each utility chapter.

Running utilities on objects with the DEFINE NO attribute

With DB2 version 7, you can run certain online utilities on table spaces or index spaces that were defined with the DEFINE NO attribute. In this situation, the physical creation of the data set is deferred until data is first inserted in the table space.

You can populate table spaces with undefined data sets by using the LOAD utility with either the RESUME keyword, the REPLACE keyword, or both. Using LOAD in this manner results in the following actions:

1. DB2 allocates the data sets.
2. DB2 updates the SPACE column in the catalog table to show that data sets exist.
3. DB2 loads the specified table space.

For a partitioned table space, all partitions are allocated even if the LOAD is loading only one part. Avoid attempting to populate a partitioned table space with concurrent LOAD PART jobs until after one of the jobs has caused all the data sets to be created.

The following online utilities will process a table space that has been defined but has indexes that remain undefined:

- CHECK DATA
- CHECK INDEX
- COPY TABLESPACE
- MERGECOPY
- MODIFY RECOVERY
- QUIESCE WRITE(NO)
- REBUILD INDEX
- RECOVER
- REORG TABLESPACE
- REPAIR, except REPAIR DBD

- RUNSTATS TABLESPACE INDEX(ALL)
- UNLOAD

Online utilities that encounter an undefined target object may issue informational message DSNU185I but continue processing.

No stand-alone utilities can be used on objects where the data sets have not been defined.

Part 2. DB2 online utilities

Chapter 3. Invoking DB2 online utilities . . . 23
Creating utility control statements . . . 23
Control statement coding rules . . . 23
Example of option description . . . 24
Data sets used by online utilities . . . 24
Concatenating data sets . . . 27
Controlling data set disposition . . . 27
Security . . . 28
Using the DB2 utilities panel in DB2I . . . 28
Using the DSNU CLIST command in TSO . . . 31
DSNU CLIST command syntax . . . 32
DSNU CLIST option descriptions . . . 32
Reviewing DSNU CLIST command output . . . 36
Editing the generated JCL data set . . . 37
Examples . . . 37
Using the supplied JCL procedure (DSNUPROC) . . . 38
DSNUPROC syntax . . . 38
DSNUPROC option descriptions . . . 38
Sample DSNUPROC listing . . . 40
Creating the JCL data set yourself . . . 41
EXEC statement . . . 41

Chapter 4. Monitoring and controlling online utilities . . . 43
Monitoring utilities with the DISPLAY UTILITY command . . . 43
Determining the status of a utility . . . 43
Determining which utility phase is currently executing . . . 44
Determining why a utility failed to complete . . . 44
Running utilities concurrently . . . 44
Running online utilities in a data sharing environment . . . 45
Terminating an online utility with the TERM UTILITY command . . . 45
Restarting an online utility . . . 46
Updating the JCL data set for restarting a utility . . . 46
Adding or deleting utility statements . . . 47
Restarting after the output data set is full . . . 47
RESTARTING with TEMPLATES . . . 47
RESTARTING with LISTS . . . 48
Other restart hints . . . 48

Chapter 5. CATMAINT . . . 49
Syntax and options of the control statement . . . 49
Syntax diagram . . . 49
Option descriptions . . . 49
Instructions for running CATMAINT . . . 49
Before running CATMAINT . . . 50
Data sets used by CATMAINT . . . 50
Instructions for specific tasks . . . 50
Terminating or restarting CATMAINT . . . 50
Concurrency and compatibility . . . 51

Chapter 6. CHECK DATA . . . 53
Syntax and options of the control statement . . . 54
Syntax diagram . . . 54
Option descriptions . . . 54

© Copyright IBM Corp. 1983, 2001

Instructions for running CHECK DATA . . . 58
Before running CHECK DATA . . . 58
For a table with no LOB columns: . . . 58
For a table with LOB columns: . . . 58
Create exception tables: . . . 58
Exception processing a table with a LOB column: . . . 59
Example: creating an exception table for the project activity table . . . 60
Complete all LOB column definitions . . . 60
Data sets used by CHECK DATA . . . 61
Creating the control statement . . . 62
Instructions for specific tasks . . . 62
Specify the scope of CHECK DATA . . . 62
Checking several table spaces . . . 62
Finding violations . . . 62
Detecting and correcting constraint violations . . . 62
Resetting CHECK-pending status . . . 63
Interpreting LOB column errors . . . 63
Resetting auxiliary CHECK-pending status . . . 64
Terminating or restarting CHECK DATA . . . 64
Terminating CHECK DATA . . . 64
Restarting CHECK DATA . . . 64
Concurrency and compatibility . . . 65
Sample control statements . . . 66

Chapter 7. CHECK INDEX . . . 67
Syntax and options of the control statement . . . 67
Syntax diagram . . . 68
Option descriptions . . . 68
Instructions for running CHECK INDEX . . . 69
Data sets used by CHECK INDEX . . . 70
Creating the control statement . . . 71
Instructions for specific tasks . . . 71
Checking a single logical partition . . . 71
Reviewing CHECK INDEX output . . . 71
Terminating or restarting CHECK INDEX . . . 72
Concurrency and compatibility . . . 72
Sample control statements . . . 73

Chapter 8. CHECK LOB . . . 75
Syntax and options of the control statement . . . 75
Syntax diagram . . . 76
Option descriptions . . . 76
Instructions for running CHECK LOB . . . 77
Before running CHECK LOB . . . 78
Data sets used by CHECK LOB . . . 78
Creating the control statement . . . 78
Instructions for specific tasks . . . 78
Finding and resolving violations . . . 79
Resetting CHECK-pending status for a LOB table space . . . 79
Resolving media failure . . . 79
Terminating or restarting CHECK LOB . . . 79
Terminating CHECK LOB . . . 79
Restarting CHECK LOB . . . 79
Concurrency and compatibility . . . 79
Sample control statements . . . 80

Chapter 9. COPY 81
Syntax and options of the control statement 82
Syntax diagram 82
Option descriptions 84
Instructions for running COPY 89
Before running COPY 90
Data sets used by COPY 90
Creating the control statement 92
Instructions for specific tasks 92
Making full image copies 92
Making incremental image copies 93
Making multiple image copies 93
Copying partitions or data sets in separate jobs 95
Copying a list of objects 95
Using more than one COPY statement 97
Copying segmented table spaces 97
Using DFSMS concurrent copy 97
Specifying conditional image copies 98
Preparing for recovery 99
Improving performance 100
Considerations for running COPY 100
Copying table spaces with mixed volume IDs 100
Defining generation data groups 101
Using DB2 with DFSMS products 101
Putting image copies on tape 101
Copying a LOB table space 101
Terminating or restarting COPY 101
Warning against TERM UTILITY 101
Implications of DISP on the DD statement 102
Restarting with a new data set 102
Restarting a COPY job 102
Restarting COPY after an out of space condition 102
Concurrency and compatibility 102
Sample control statements 104

Chapter 10. COPYTOCOPY 109
Syntax and options of the control statement 110
Syntax diagram 110
Option descriptions 112
Instructions for running COPYTOCOPY 115
Before running COPYTOCOPY 115
Data sets used by COPYTOCOPY 115
Creating the control statement 116
Instructions for specific tasks 116
Copying full or incremental image copies 117
Copying incremental image copies 117
Using more than one COPYTOCOPY statement 117
Copying inline copy made by REORG with part range 117
Copying from a specific image copy 117
Using TEMPLATE with COPYTOCOPY 118
Considerations for running COPYTOCOPY 118
Updating SYSIBM.SYSCOPY records 118
Determining the input copy used 118
Retaining tape mounts 118
Defining generation data groups 119
Using DB2 with DFSMS products 119

Putting image copies on tape . . . 119
Copying a LOB table space . . . 119
Terminating or restarting COPYTOCOPY . . . 119
Restarting a COPYTOCOPY job . . . 119
Restarting COPYTOCOPY after an out of space condition . . . 119
Concurrency and compatibility . . . 119
Sample control statement . . . 120

Chapter 11. DIAGNOSE . . . 123
Syntax and options of the control statement . . . 123
Syntax diagram . . . 123
Option descriptions . . . 125
Instructions for running DIAGNOSE . . . 127
Data sets used by DIAGNOSE . . . 127
Instructions for specific tasks . . . 128
Forcing a utility abend with DIAGNOSE . . . 128
Terminating or restarting DIAGNOSE . . . 128
Concurrency and compatibility . . . 128
Sample control statements . . . 128

Chapter 12. EXEC SQL . . . 131
Syntax and options of the control statement . . . 131
Syntax diagram . . . 131
Option descriptions . . . 131
Concurrency and compatibility . . . 132
Sample control statements . . . 132

Chapter 13. LISTDEF . . . 135
Syntax and options of the control statement . . . 136
Syntax diagram . . . 136
Option descriptions . . . 138
Instructions for using LISTDEF . . . 143
Generating object lists using LISTDEF . . . 144
Valid pattern-matching characters . . . 145
How to use LISTDEF specifications . . . 146
Key LISTDEF operations . . . 146
Specifying catalog and directory objects in LISTDEFs . . . 147
Concurrency and compatibility . . . 147
Sample control statements . . . 147

Chapter 14. LOAD . . . 151
Syntax and options of the control statement . . . 152
Syntax diagram . . . 153
Option descriptions . . . 155
Into-table-spec . . . 167
Option descriptions for INTO TABLE . . . 169
Instructions for running LOAD . . . 182
Before running LOAD . . . 183
Loading data using a cursor . . . 183
Data sets used by LOAD . . . 183
Instructions for specific tasks . . . 186
Loading variable-length data . . . 186
Ordering loaded records . . . 186
Replacing data with LOAD . . . 187
Using LOAD for tables with identity columns or ROWID columns . . . 188
Adding more data to a table or partition . . . 189

Deleting all the data in a table space . . . 189
Loading partitions . . . 189
Loading data with referential constraints . . . 190
Correcting referential constraint violations . . . 192
Compressing data . . . 192
Loading data from DL/I . . . 193
Loading data with a dynamic SQL statement . . . 194
Using inline COPY with LOAD . . . 194
Improving performance . . . 194
Improving performance for parallel processing . . . 195
Improving performance with SORTKEYS . . . 196
Improving performance with LOAD or REORG PREFORMAT . . . 197
Considerations for running LOAD . . . 198
Converting input data . . . 198
Specifying input fields . . . 199
Building indexes while loading data . . . 200
Building indexes in parallel for LOAD . . . 200
Leaving free space . . . 202
Loading with RECOVER-pending or REBUILD-pending status . . . 202
Using exit procedures . . . 203
Loading columns defined as ROWID . . . 203
Loading a LOB column . . . 204
Using LOAD LOG on a LOB table space . . . 204
Inline statistics collection for discarded rows . . . 204
Inline COPY for a base table space . . . 204
Terminating or restarting LOAD . . . 204
Restarting after an out of space condition . . . 206
Concurrency and compatibility . . . 206
After running LOAD . . . 208
Copying the loaded table space or partition . . . 209
Resetting the COPY-pending status . . . 209
Resetting the REBUILD-pending status . . . 209
Resetting the CHECK-pending status . . . 210
DELETE YES . . . 210
Exception tables . . . 210
Error and sort data sets . . . 211
SCOPE PENDING . . . 212
Collecting inline statistics while loading a table . . . 212
Running CHECK INDEX after loading a table having indexes . . . 212
Recovering a failed LOAD job . . . 212
Reorganizing an auxiliary index after LOAD . . . 212
Sample control statements . . . 213

Chapter 15. MERGECOPY . . . 221
Syntax and options of the control statement . . . 222
Syntax diagram . . . 222
Option descriptions . . . 222
Instructions for running MERGECOPY . . . 224
Data sets used by MERGECOPY . . . 225
Creating the control statement . . . 226
Instructions for specific tasks . . . 226
Specifying full or incremental image copy . . . 226
Merging inline copies . . . 227
Using MERGECOPY with individual data sets . . . 227
Deciding between MERGECOPY or COPY . . . 227
Avoiding MERGECOPY LOG RBA inconsistencies . . . 227

Creating image copies in a JES3 environment . . . 228
Running MERGECOPY on the directory . . . 228
Terminating or restarting MERGECOPY . . . 228
Concurrency and compatibility . . . 228
Sample control statements . . . 229

Chapter 16. MODIFY RECOVERY . . . 231
Syntax and options of the control statement . . . 232
Syntax diagram . . . 232
Option descriptions . . . 232
Instructions for running MODIFY RECOVERY . . . 233
Before running MODIFY RECOVERY . . . 234
Data sets used by MODIFY RECOVERY . . . 234
Creating the control statement . . . 234
Instructions for specific tasks . . . 234
Deleting SYSLGRNX rows . . . 234
Deleting all image copy entries . . . 235
Deleting recovery rows for indexes . . . 235
Reclaiming space in the DBD . . . 235
Improving REORG performance after adding a column . . . 235
Terminating or restarting MODIFY RECOVERY . . . 235
Concurrency and compatibility . . . 235
Sample control statements . . . 236

Chapter 17. MODIFY STATISTICS . . . 237
Syntax and options of the control statement . . . 237
Syntax diagram . . . 238
Option descriptions . . . 238
Instructions for running MODIFY STATISTICS . . . 239
Data sets used by MODIFY STATISTICS . . . 240
Creating the control statement . . . 240
Instructions for specific tasks . . . 240
Deciding which statistics history rows to delete . . . 240
Deleting specific statistics history rows . . . 241
Terminating or restarting MODIFY STATISTICS . . . 241
Concurrency and compatibility . . . 241
Sample control statements . . . 241

Chapter 18. OPTIONS . . . 245
Syntax and options of the control statement . . . 245
Syntax diagram . . . 245
Option descriptions . . . 246
Instructions for using OPTIONS . . . 248
Concurrency and compatibility . . . 249
Sample control statements . . . 249

Chapter 19. QUIESCE . . . 251
Syntax and options of the control statement . . . 251
Syntax diagram . . . 251
Option descriptions . . . 252
Instructions for running QUIESCE . . . 253
Before running QUIESCE . . . 253
Data sets used by QUIESCE . . . 253
Creating the control statement . . . 254
Instructions for specific tasks . . . 254
Using QUIESCE for recovery . . . 254

Using QUIESCE on catalog and directory objects 254
Obtaining a common quiesce point 254
Specifying a list of table spaces and table space sets 255
Considerations for running QUIESCE 255
Terminating or restarting QUIESCE 256
Concurrency and compatibility 256
Sample control statements 258

Chapter 20. REBUILD INDEX 259
Syntax and options of the control statement 259
Syntax diagram 259
Option descriptions 260
Instructions for running REBUILD INDEX 264
Before running REBUILD INDEX 264
Data sets used by REBUILD INDEX 264
Creating the control statement 265
Instructions for specific tasks 265
Rebuilding index partitions 266
Improving performance and space utilization 266
Building indexes in parallel for REBUILD INDEX 266
Resetting the REBUILD-pending status 269
Rebuilding critical catalog indexes 270
Rebuilt index recoverability 270
Terminating or restarting REBUILD INDEX 270
Concurrency and compatibility 271
Sample control statements 272

Chapter 21. RECOVER 275
Syntax and options of the control statement 276
Syntax diagram 276
Option descriptions 277
Instructions for running RECOVER 282
Before running RECOVER 282
Data sets used by RECOVER 283
Instructions for specific tasks 283
Recovering a table space 283
Recovering a list of objects 284
Recovering a data set or partition 284
Recovering with incremental copies 285
Recovering a page 285
Recovering an error range 286
Recovering with a data set copy not made by DB2 286
Recovering catalog and directory objects 287
Recovering a table space containing LOB data 290
Performing a point-in-time recovery 291
Avoiding specific image copy data sets 295
Improving performance 295
Optimizing the LOGAPPLY phase 295
Recovering image copies in a JES3 environment 297
Resetting RECOVER-pending or REBUILD-pending status 297
Considerations for running RECOVER 297
Allocating incremental image copies 297
Performing fallback recovery 298
Retaining tape mounts 298
Avoiding damaged media 299
Recovering table spaces and index spaces with mixed volume IDs 299

Terminating or restarting RECOVER 299
Concurrency and compatibility 300
Sample control statements 301

Chapter 22. REORG INDEX 303
Syntax and options of the control statement 304
Syntax diagram 304
Option descriptions 306
Instructions for running REORG INDEX 314
Before running REORG INDEX 314
Data sets used by REORG INDEX 316
Creating the control statement 316
Instructions for specific tasks 316
Determining when an object should be reorganized 317
Determining when an index requires reorganization 317
Specifying access with SHRLEVEL 318
Unloading without reloading. 319
Considerations for fallback recovery. 319
Changing data set definitions 319
Temporarily interrupting REORG 320
Improving performance 320
Terminating or restarting REORG INDEX 321
Concurrency and compatibility 322
REORG INDEX compatibility 323
Reviewing REORG INDEX output 324
Sample control statements 325

Chapter 23. REORG TABLESPACE 327
Syntax and options of the control statement 329
Syntax diagram 329
Option descriptions 335
REORG TABLESPACE options syntax 355
Option descriptions for REORG TABLESPACE options 355
Instructions for running REORG TABLESPACE 356
Before running REORG TABLESPACE 357
Data sets used by REORG TABLESPACE 361
Creating the control statement 364
Instructions for specific tasks 365
Determining when an object should be reorganized 365
Specifying access with SHRLEVEL 367
Omitting the output data set. 368
Unloading without reloading. 368
Reclaiming space from dropped tables. 369
Considerations for fallback recovery. 369
Reorganizing the catalog and directory 369
Changing data set definitions 371
Temporarily interrupting REORG 371
Building a compression dictionary 372
Overriding dynamic DFSORT and SORTDATA allocation 372
Rebalancing partitions using REORG 372
Using inline COPY with REORG TABLESPACE 373
Improving performance 373
Considerations for running REORG 378
Sorting data in clustering order 378
Methods of unloading data 378
Encountering an error in the RELOAD phase 379

Reorganizing partitioned table spaces . . . 379
Reorganizing segmented table spaces . . . 379
Counting records loaded during RELOAD phase . . . 379
Reorganizing a LOB table space . . . 380
Terminating or restarting REORG TABLESPACE . . . 380
Concurrency and compatibility . . . 383
REORG TABLESPACE compatibility . . . 383
Reviewing REORG TABLESPACE output . . . 387
After running REORG TABLESPACE . . . 387
Sample control statements . . . 388

Chapter 24. REPAIR . . . 397
Syntax and options of the control statement . . . 398
Syntax diagram . . . 398
REPAIR option descriptions . . . 398
SET TABLESPACE and SET INDEX statement syntax . . . 399
SET TABLESPACE and SET INDEX option descriptions . . . 400
LOCATE block syntax . . . 401
LOCATE TABLESPACE statement option descriptions . . . 401
LOCATE INDEX statement option descriptions . . . 403
VERIFY statement syntax . . . 403
VERIFY statement option descriptions . . . 404
REPLACE statement syntax . . . 404
REPLACE statement option descriptions . . . 405
DELETE statement syntax and description . . . 405
DUMP statement syntax . . . 406
DUMP statement option descriptions . . . 406
DBD statement syntax . . . 407
DBD statement option descriptions . . . 408
Instructions for running REPAIR . . . 409
Before running REPAIR . . . 409
Making a copy of the table space . . . 409
Restoring damaged indexes . . . 410
Data sets used by REPAIR . . . 410
Creating the control statement . . . 410
Instructions for specific tasks . . . 410
Resetting table space status . . . 411
Resetting index space status . . . 411
Repairing a damaged page . . . 411
Using the DBD statement . . . 411
Locating rows by key . . . 412
Using VERIFY, REPLACE, and DELETE operations . . . 412
Repairing critical catalog table spaces and indexes . . . 412
Terminating or restarting REPAIR . . . 413
Concurrency and compatibility . . . 413
Reviewing REPAIR output . . . 415
After running REPAIR . . . 416
Sample control statements . . . 416

Chapter 25. REPORT . . . 419
Syntax and options of the control statement . . . 420
Syntax diagram . . . 420
Option descriptions . . . 420
Instructions for running REPORT . . . 423
Data sets used by REPORT . . . 423
Creating the control statement . . . 424

Instructions for specific tasks 424
Reporting recovery information 424
Running REPORT on the catalog and directory 425
Terminating or restarting REPORT 425
Concurrency and compatibility 425
Reviewing REPORT output 426
Sample control statements 429

Chapter 26. RUNSTATS 431
Syntax and options of the control statement 432
RUNSTATS TABLESPACE syntax diagram 432
RUNSTATS TABLESPACE option descriptions 434
RUNSTATS INDEX syntax diagram 438
RUNSTATS INDEX option descriptions 438
Instructions for running RUNSTATS 441
Before running RUNSTATS 441
Data sets used by RUNSTATS 442
Creating the control statement 442
Instructions for specific tasks 442
Deciding when to use RUNSTATS 442
Assessing table space status 442
Updating statistics for a partitioned table space 442
Running RUNSTATS on the DB2 catalog 443
Improving performance 443
Terminating or restarting RUNSTATS 443
Concurrency and compatibility 444
Reviewing RUNSTATS output 445
Access path statistics 447
Space statistics (columns for tuning information) 449
After running RUNSTATS 453
Sample control statements 453

Chapter 27. STOSPACE 457
Syntax and options of the control statement 457
Syntax diagram 457
Option descriptions 457
Instructions for running STOSPACE 458
Data sets used by STOSPACE 458
Creating the control statement 458
Instructions for specific tasks 459
Ensuring availability of objects required by STOSPACE 459
Obtaining statistical information with STOSPACE 459
Understanding the values in a SPACE column 460
Considerations for running STOSPACE 460
Terminating or restarting STOSPACE 460
Concurrency and compatibility 460
Reviewing STOSPACE output 461
Sample control statement 461

Chapter 28. TEMPLATE 463
Syntax and options of the control statement 463
Syntax diagram 463
Option descriptions 464
Instructions for using TEMPLATE 472
Key TEMPLATE operations 473
Creating data set names 473

Controlling data set size 474
Default space calculations 474
Working with TAPE 474
Working with GDGs. 475
Concurrency and compatibility 475
Sample control statements 475

Chapter 29. UNLOAD 477
Syntax and options of the control statement 477
Syntax diagram 477
Option descriptions 478
FROM-TABLE-spec. 485
Option descriptions for FROM TABLE 489
Instructions for running UNLOAD 507
Data sets used by UNLOAD 508
Instructions for specific tasks 508
Unloading partitions. 509
Selecting tables and rows to unload. 509
Selecting and ordering columns to unload 509
Using inline image copies with UNLOAD 509
Converting data with the UNLOAD utility 511
Specifying output field types 511
Specifying output field positioning and size 513
Determining the layout of output fields 514
Specifying TRUNCATE and STRIP options for output data 516
Generating LOAD statements 517
Unloading compressed data 517
Interpreting field specification errors 518
Terminating or restarting UNLOAD 518
Concurrency and compatibility 518
Sample control statements 520

Chapter 3. Invoking DB2 online utilities

This chapter contains procedures and guidelines for creating utility control statements and describes five methods for invoking the DB2 utilities.

Creating utility control statements is the first step required to run an online utility.

After creating the utility statements, use one of the following methods for invoking the online utilities:

- "Using the DB2 utilities panel in DB2I" on page 28
- "Using the DSNU CLIST command in TSO" on page 31
- "Using the supplied JCL procedure (DSNUPROC)" on page 38
- "Creating the JCL data set yourself" on page 41
- "Invoking utilities as a stored procedure (DSNUTILS)" on page 631

For the least involvement with JCL, use either DB2I or the DSNU CLIST command, and then edit the generated JCL to alter or add necessary fields on the JOB or ROUTE cards before submitting the job. Both of these methods require TSO, and the first also requires access to the DB2 Utilities Panel in DB2 Interactive (DB2I).

If you want to work with JCL or create your own JCL, choose the third or fourth method listed previously.

Creating utility control statements

Utility control statements define the function that the utility job performs.

You can create the utility control statements with the ISPF/PDF edit function. After they are created, save them in a sequential or partitioned data set.

Control statement coding rules

Utility control statements are typically read from the SYSIN data set. LISTDEF control statements may be read from the SYSLISTD data set. TEMPLATE control statements may be read from the SYSTEMPL data set. The statements in these data sets must obey these rules:

- If the records are 80-character fixed-length records, columns 73 through 80 are ignored.
- The records are concatenated before being parsed; therefore, a statement or any of its syntactical constructs can span more than one record. No continuation character is necessary.
- The control statements must begin with one of these *online utility names*:

CATMAINT	LOAD	REORG INDEX
CHECK DATA	MERGECOPY	REORG TABLESPACE
CHECK INDEX	MODIFY RECOVERY	REPAIR
CHECK LOB	MODIFY STATISTICS	REPORT
COPY	OPTIONS	RUNSTATS
DIAGNOSE	QUIESCE	STOSPACE
LISTDEF	REBUILD INDEX	TEMPLATE
	RECOVER	UNLOAD

  At least one blank character must follow the name.

© Copyright IBM Corp. 1983, 2001

- Other syntactical constructs in the utility control statement describe options; you can separate these constructs with an arbitrary number of blanks.
- The SYSIN stream can contain multiple utility control statements.

The online utility name determines which options can follow it. You can specify more than one utility control statement in the SYSIN stream. Be aware that, if any of the control statements returns a return code of 8 or greater, the subsequent statements in the job step are not executed.

Options are typically described by an *option keyword*, followed by a *parameter.* The parameter value can be a keyword. Values of parameters are sometimes enclosed in parentheses. The syntax diagrams for utility control statements included in this chapter show parentheses where they are required.

You can enter comments within the SYSIN stream. Comments must begin with two hyphens (**--**) and are subject to the following rules:

- The two hyphens must be on the same line, not separated by a space
- You can start comments wherever a space is valid, except within a delimiter token.
- Comments are terminated by the end of the line, for example:

```
// SYSIN DD *
   RUNSTATS TABLESPACE DSNDB06.SYSDBASE  -- COMMENT HERE
   -- COMMENT HERE
/*
```

Example of option description

Where the syntax of each utility control statement is described, parameters are indented under the option keyword they must follow. Here is a typical example:

WORKDDN *ddname*

Specifies a temporary work file.

ddname is the data set name of the temporary file.

The **default** is **SYSUT1**.

In the example, WORKDDN is an option keyword, and *ddname* is a parameter. As noted previously, you can enclose parameter values in parentheses but they are not always required. The description of the temporary work file can be written as either `WORKDDN SYSUT1` or `WORKDDN (SYSUT1)`.

Data sets used by online utilities

Every online utility job requires a SYSIN DD statement to describe an input data set; some utilities also require other data sets. Table 1 on page 25 lists the name of each DD statement that might be needed, the online utilities that require it, and the purpose of the corresponding data sets. If an alternate DD statement is allowed, you specify it as a parameter in a utility option. Table 1 on page 25 also lists the option keywords that you can use. DCB attributes that you specify on the DD statement are referred to as *user-specified values.*

Table 1. Data sets used by online utilities

DD Name	Used By	Purpose	Option Keyword
ddname	COPY[8]	A single data set DB2 uses when you specify the FILTERDDN option in the utility control statement; contains a list of VSAM data set names used during COPY using the CONCURRENT and FILTERDDN options.	FILTERDDN
SORTDATA	REORG	Work data set for sorting data, where *nn* is a 2-digit number. You can use several data sets. To estimate the size of the data set needed, see "Data sets used by REORG TABLESPACE" on page 361.	SHRLEVEL CHANGE
DSSPRINT	COPY	Output data set for messages; required when CONCURRENT copy is used and the SYSPRINT DD card points to a data set.	CONCURRENT
SORTOUT	CHECK DATA[2,7], LOAD[1,3,7], REORG[1,7]	Holds sorted keys (sort output) and allows the SORT phase to be restarted; for CHECK DATA, holds sorted keys (sort output).	WORKDDN
SORTWK*nn* [4]	CHECK DATA, CHECK INDEX, LOAD, REBUILD INDEX, REORG	Work data set for sorting indexes where *nn* is a 2-digit number. You can use several data sets. To estimate the size of the data set required, see page 61 for CHECK DATA, 70 for CHECK INDEX, 184 for LOAD, 265 for REBUILD INDEX, or 364 for REORG.	None
SW*mm*WK*nn* [4]	LOAD, REBUILD INDEX, REORG	Optional work data sets for sorting index keys using the SORTKEYS keyword, where *mm* and *nn* are 2-digit numbers. You can use several data sets. To estimate the size of the data set required, see page 202 for LOAD, 269 for REBUILD INDEX, or 377 for REORG.	None
SYSCOPY	COPY, MERGECOPY, LOAD[5], REORG [5]	An output data set for copies.	COPYDDN RECOVERYDDN
SYSDISC	LOAD, REORG DISCARD, optional for REORG	Contains discarded records (optional).	DISCARDDN
SYSERR	CHECK DATA[2], LOAD	Contains information about errors encountered during processing.	ERRDDN
SYSIN	All utilities	An input data set for utility statements.	None
SYSMAP	LOAD[3]	Contains information about what input records violated a constraint.	MAPDDN

Table 1. Data sets used by online utilities (continued)

DD Name	Used By	Purpose	Option Keyword
SYSPRINT	All utilities	Messages and printed output (usually SYSOUT).	None
SYSPUNCH	REORG, UNLOAD	The utility contains a LOAD statement that is generated by REORG or UNLOAD. For REORG, the LOAD statement loads records that REORG DISCARD or REORG UNLOAD EXTERNAL wrote to the DISCARD or UNLOAD data sets.	PUNCHDDN
SYSREC	LOAD[2], REORG[6], UNLOAD[2]	The utility contains the input data set of LOAD, unloaded records for REORG, or UNLOAD.	INDDN UNLDDN
SYSUT1	CATMAINT, CHECK DATA[7], CHECK INDEX[2], LOAD[1,3,7], MERGECOPY, REBUILD INDEX, REORG[1,7]	A temporary work data set that holds sorted keys for input to the SORT phase; for MERGECOPY, it holds intermediate merged output.	WORKDDN
UTPRIN*mm*	LOAD, REBUILD INDEX, REORG	Optional print message data sets, used when the SORTKEYS keyword is specified, where *mm* is a 2-digit number.	None
UTPRINT	CHECK DATA, CHECK INDEX, LOAD, REORG, REBUILD INDEX	Contains messages from DFSORT™ (usually, SYSOUT or DUMMY).	None

Notes:

[1] Required for tables with indexes, not required when using REORG with the SORTKEYS option.

[2] Required.

[3] When referential constraints exist and ENFORCE(CONSTRAINTS) is specified.

[4] If tape is specified, the maximum key length of all indexes involved in the sort phase must be a minimum of 6 bytes. This length, when added to the internally assigned 12-byte header, must be at least 18 bytes as required by DFSORT.

[5] Required for LOAD with COPYDDN or RECOVERYDDN and for REORG with COPYDDN, RECOVERYDDN, SHRLEVEL REFERENCE, or SHRLEVEL CHANGE.

[6] Required unless you specify NOSYSREC or SHRLEVEL CHANGE.

[7] Data sets can not be shared between SORTOUT and SYSUT1. Sharing these data sets can cause unpredictable results.

[8] If you specify FILTERDDN, there is no default DD name. You must supply a name.

For input data sets, the online utilities use the logical record length (LRECL), record format (RECFM), and block size (BLKSIZE) with which the data set was created. Variable spanned (VS) or variable blocked spanned (VBS) record formats are not allowed for utility input data sets. The only exception is for the LOAD utility, which accepts unloaded SQL/DS™ data in VBS format.

For output data sets, the online utilities determine both the logical record length and the record format. If you supply values for LRECL or RECFM, they are ignored.

If you supply block size, it is used; otherwise, the utility chooses a block size appropriate for the storage device. Partitioned data sets (PDS) are not allowed for output data sets.

For both input and output data sets, the online utilities use the value you supply for the number of buffers (BUFNO), with a maximum of 99 buffers. The default number of buffers is 20. The utilities set the number of channel programs equal to the number of buffers. The parameters used to specify the buffer size (BUFSIZE) and the number of channel programs (NCP) are ignored. If you omit any DCB parameters, the utilities choose default values.

Note: Increasing the number of buffers (BUFNO) can result in an increase in real
storage utilitzation and page fixing below 16MB.

The default for the FASTSWITCH parameter is YES. If you do not specify
FASTSWITCH NO, the fifth level qualifier in the data set name alternates between
I0001 and J0001.

Restriction: DB2 does not support the undefined record format (RECFM=U) for any data set.

Concatenating data sets

DB2 utilities let you concatenate unlike input data sets. Therefore, the data sets in a concatenation list can have differing block sizes, logical record lengths, and record formats. If you want to concatenate variable and fixed blocked data sets, the logical record length must be eight bytes smaller than the block size.

You cannot concatenate output data sets.

Controlling data set disposition

Most data sets need to exist only during utility execution (for example, during reorganization). However, you must keep several data sets in certain circumstances:

- Retain the image copy data sets until they are no longer needed for recovery.
- Retain the unload data sets if you specify UNLOAD PAUSE, UNLOAD ONLY, UNLOAD EXTERNAL or DISCARD for the REORG utility.
- Retain the SYSPUNCH data set if you specify UNLOAD EXTERNAL or DISCARD for the REORG utility until the contents are no longer needed for subsequent loads.
- Retain the discard data set until the contents are no longer needed for subsequent loads.

Because you might need to restart a utility, take the following precautions when defining the disposition of data sets:

| - Use DISP=(NEW,CATLG,CATLG) or DISP=(MOD,CATLG) for data sets you want
| to retain.
- Use DISP=(MOD,DELETE,CATLG) for data sets that you want to discard after utility execution.
- Use DISP=(NEW,DELETE) for DFSORT SORTWKnn data sets, or refer to *DFSORT Application Programming: Guide* for alternatives.
- Do not use temporary data set names.

Refer to Table 113 on page 632 and Table 114 on page 633 for information about the default data dispositions specified for dynamically-allocated data sets.

Security

To prevent unauthorized access to data sets (for example, image copies), you can protect the data sets with the Resource Access Control Facility (RACF). To use a utility with a data set protected by RACF, you must be authorized to access the data set.

Using the DB2 utilities panel in DB2I

Using the DB2 Utilities panel to execute DB2 online utilities requires the least knowledge of JCL.

Restriction for using the DB2 Utilities panel: You cannot use the DB2 Utilities panel in DB2I to submit a COPY job for a list of objects (with or without the CONCURRENT keyword).

Editing and submitting a utility job: If your site does not have default JOB and ROUTE statements, you must edit the JCL to define them. If you edit the utility job before submitting it, you must use the ISPF editor and submit your job directly from the editor. Use the procedure outlined in the following example:

1. Create the utility control statement for the online utility you intend to execute, and save it in a sequential or partitioned data set.

   For example, the following utility control statement makes an incremental image copy of table space DSN8D71A.DSN8S71D with a SHRLEVEL value of CHANGE:

   ```
   COPY TABLESPACE DSN8D71A.DSN8S71D
     FULL NO
     SHRLEVEL CHANGE
   ```

   For the rest of this example, suppose that you save the statement in the default data set, UTIL.
2. Select the DB2I menu from the ISPF Primary Option Menu.
3. Select the UTILITIES option on the DB2I Utilities panel. Items you must specify are highlighted on the DB2 Utilities panel in Figure 1.

```
DSNEUP01                        DB2 UTILITIES
 ===>

 Select from the following:

  1 FUNCTION ===> EDITJCL   (SUBMIT job, EDITJCL, DISPLAY, TERMINATE)
  2 JOB ID   ===> TEMP      (A unique job identifier string)
  3 UTILITY  ===> COPY      (CHECK DATA, CHECK INDEX, CHECK LOB,
                             COPY, DIAGNOSE, LOAD, MERGE, MODIFY,
                             QUIESCE, REBUILD, RECOVER, REORG INDEX,
                             REORG LOB, REORG TABLESPACE, REPORT,
                             REPAIR, RUNSTATS, STOSPACE, UNLOAD.)

  4 STATEMENT DATA SET    ===> UTIL

 Specify restart or preview option, otherwise specify NO.

  5 RESTART  ===> NO         (NO, CURRENT, PHASE or PREVIEW)

  6 LISTDEF? (YES|NO) ===>          TEMPLATE? (YES|NO) ===>

  *  The data set names panel will be displayed when required by a utility.


 PRESS:  ENTER to process     END to exit    HELP for more information
```

Figure 1. DB2 utilities panel

4. Fill in field 1 with the function you want to execute. In this example, you want to submit the utility job, but you want to edit the JCL first. After you have edited the utility job, specify SUBMIT on the editor command line.
5. Field 2 must be a unique identifier for your utility job. The default value TEMP. In this example, that value is satisfactory; leave it as is.
6. Fill in field 3 with the utility you want to run. Specify REORG LOB to indicate REORG TABLESPACE of a LOB table space.

   In this example, specify COPY.
7. Fill in field 4 if you want to use an input data set other than the default data set. Unless you enclose the data set name between apostrophes, TSO adds your user identifier as a prefix. In this example, specify UTIL, which is the default data set.
8. Change field 5 if this job restarts a stopped utility or to execute a utility in PREVIEW mode. In this example, leave the default value, NO.
9. Specify in field 6 whether you are using LISTDEF statements or TEMPLATE statements in this utility. When you use LISTDEF statements, you need to enter values in the library data set panel. DB2 displays a library data set panel when you use TEMPLATE statements, but the field entries are optional.
10. Press ENTER.

If you specify COPY, LOAD, MERGECOPY, REORG TABLESPACE, or UNLOAD as the utility in field 3, you must complete the fields on the “Data Set Names” panel. In this example, COPY was specified.

If LISTDEF Y or TEMPLATE Y is specified, DSNEUP03 panel appears.

```
DSNEUP02                          DATA SET NAMES
 ===>

   Enter data set name for LOAD or REORG TABLESPACE:
1 RECDSN   ==>

   Enter data set name for
LOAD, REORG TABLESPACE or UNLOAD:
2 DISCDSN  ==>

   Enter output data sets for local/current site for COPY, MERGECOPY,
     LOAD, or REORG:
3 COPYDSN  ==> ABC
4 COPYDSN2 ==>

   Enter output data sets for recovery site for COPY, LOAD, or REORG:
5 RCPYDSN1 ==> ABC1
6 RCPYDSN2 ==>
   Enter output data sets for REORG or UNLOAD:
7 PUNCHDSN ==>
 PRESS:  ENTER to process     END to exit    HELP for more information
```

Figure 2. Data set names panel

1. Fill in field 1 if you are running LOAD, REORG, or UNLOAD. If you are running LOAD, you must specify the data set name that contains records to be loaded. If you are running REORG or UNLOAD, you must specify the unload data set. In this example, you do not have to fill in field 1.
2. Fill in field 2 if you are running LOAD or REORG with discard processing; you must specify a discard data set. In this example, you do not have to fill in field 2.
3. Fill in field 3 with the primary output data set name for the local site if you are running COPY, LOAD, or REORG, or the current site if you are running MERGECOPY. The DDNAME generated by the panel for this field is SYSCOPY. This is an optional field for LOAD and for REORG with SHRLEVEL NONE; it is required for COPY, for MERGECOPY, and for REORG with SHRLEVEL REFERENCE or CHANGE.
4. Fill in field 4 with the backup output data set name for the local site if you are running COPY, LOAD, or REORG, or the current site if you are running MERGECOPY. The *DDNAME* generated by the panel for this field is SYSCOPY2. This is an optional field. In this example, you do not have to fill in field 4.
5. Fill in field 5 with the primary output data set for the recovery site if you are running COPY, LOAD, or REORG. The *DDNAME* generated by the panel for this field is SYSRCOPY1. This is an optional field.
6. Fill in field 6 with the backup output data set for the recovery site if you are running COPY, LOAD, or REORG. The *DDNAME* generated by the panel for this field is SYSRCOPY2. This field is optional. In this example, you do not have to fill in field 6.
7. Fill in field 7 with the output data set for the generated LOAD utility control statements if you are running REORG UNLOAD EXTERNAL, REORG DISCARD, or UNLOAD. The *DDNAME* generated by the panel for this field is SYSPUNCH. In this example, you do not have to fill in field 7.
8. Press ENTER.

This panel is presented if either LISTDEF Y or TEMPLATE Y is specified on the DSNEUP01 panel.

```
DSNEUP03         CONTROL STATEMENT DATA SET NAMES            SSID:
===>

Enter the data set name for the LISTDEF data set (SYSLISTD DD):
 1 LISTDEF DSN  ===>
    OPTIONAL or IGNORED

Enter the data set name for the TEMPLATE data set (SYSTEMPL DD):
 2 TEMPLATE DSN ===>"DSNEUV22
    OPTIONAL or IGNORED


PRESS:  ENTER to process     END to exit   HELP for more information
```

Figure 3. Control statement data set names panel

1. Fill in field 1 to specify the data set that contains a LISTDEF control statement. The default is the SYSIN data set. This field is ignored if you specified NO in field LISTDEF? in panel DSNEUP01.

   For information on using a LISTDEF control statement, see “Chapter 13. LISTDEF” on page 135.
2. Fill in field 2 to specify the data set that contains a TEMPLATE. The default is the SYSIN data set. This field is ignored if you specified NO in field TEMPLATE? in panel DSNEUP01.

   For information on using a TEMPLATE, see “Chapter 28. TEMPLATE” on page 463.

Using the DSNU CLIST command in TSO

You can also initiate a DB2 online utility by invoking the DSNU CLIST command under TSO. The CLIST command generates the JCL data set required to execute the DSNUPROC procedure and execute online utilities as batch jobs. When you use the CLIST command, you need not be concerned with details of the JCL data set.

Restriction for using the DSNU CLIST command: You cannot use the DSNU CLIST command to submit a COPY job for a list of objects (with or without the CONCURRENT keyword).

Creating a utility job: The CLIST command creates a job that performs only one utility operation. However, you can invoke the CLIST command for each utility operation you need, and then edit and merge the outputs into one job or step.

To use the DSNU CLIST command:

1. Create a file containing the required utility statements and control statements. The file is used to create the SYSIN data set in the generated job stream. Do not include double-byte character set (DBCS) data in this file.
2. Ensure that the DB2 CLIST library is allocated to the *ddname* SYSPROC.
3. Execute the command procedure by using the syntax in “DSNU CLIST command syntax” on page 32.
4. Edit the generated JCL data set to alter or add DD statements as needed.

This last step is optional. “Editing the generated JCL data set” on page 37 explains how to edit the JCL data set.

You can execute the DSNU CLIST command from the TSO command processor or the DB2I Utilities panel.

DSNU CLIST command syntax

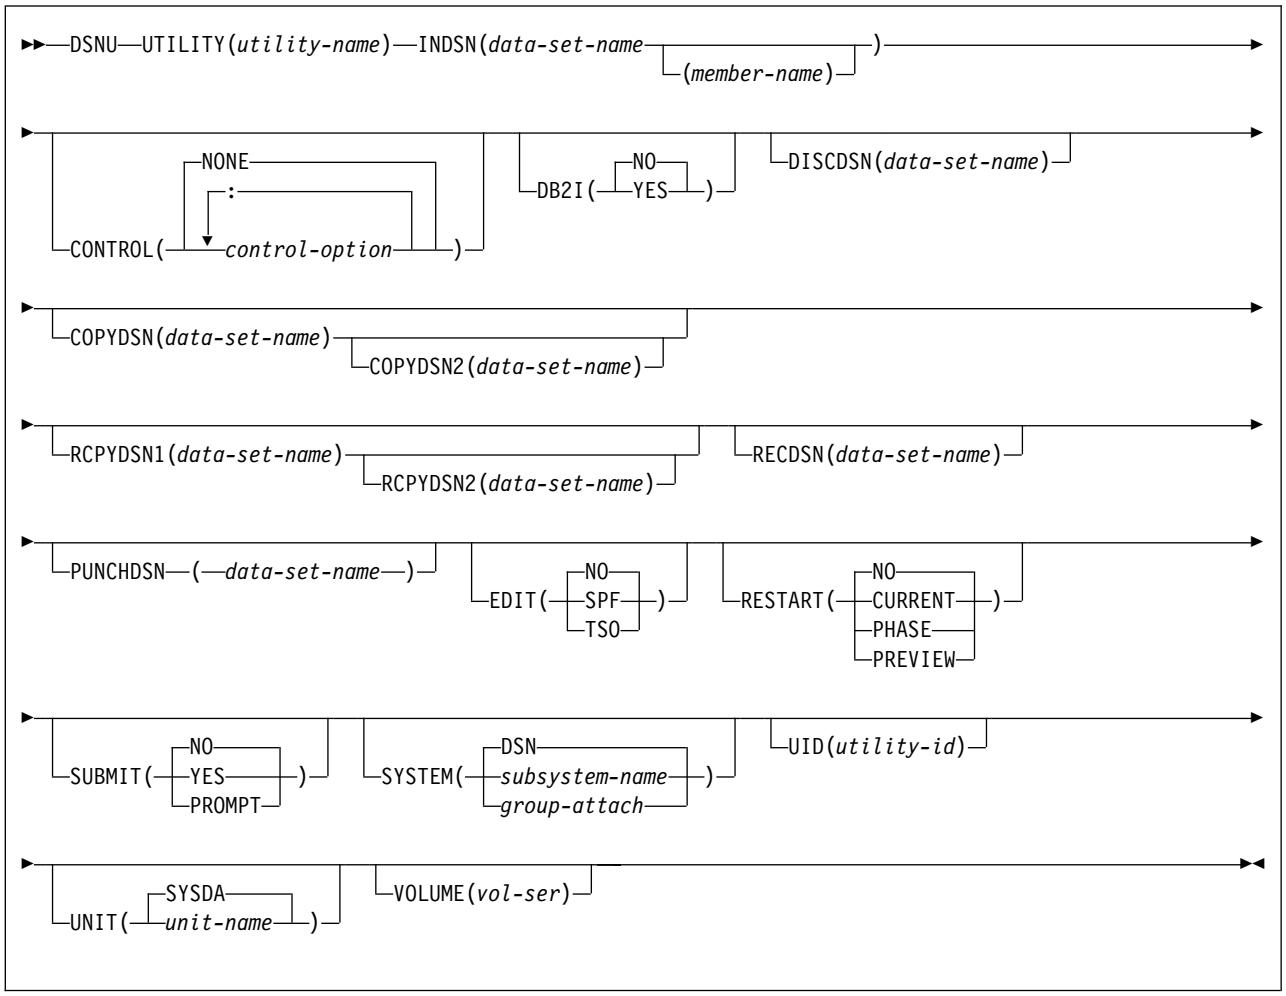

DSNU CLIST option descriptions

The parentheses shown in the following descriptions are required. If you make syntax errors or omit parameter values, TSO prompts you for the correct parameter spelling and omitted values.

% Identifies DSNU as a member of a command procedure library. Specifying this parameter is optional; however, it does improve performance.

UTILITY (*utility-name***)**
Specifies the utility you want to execute. Select the name from the following list:

CHECK DATA	MERGE	REORG LOB
CHECK INDEX	MODIFY	REORG TABLESPACE
CHECK LOB	QUIESCE	REPAIR
COPY	REBUILD	REPORT
DIAGNOSE	RECOVER	RUNSTATS
LOAD	REORG INDEX	STOSPACE
		UNLOAD

DB2 places the JCL in a data set named DSNU*xxx*.CNTL, where DSNU*xxx* is a control file name. The control file contains the statements necessary to invoke the DSNUPROC procedure which, in turn, executes the utility. If you execute another job with the same utility name, the first job is deleted. See the **UID** keyword on page 35 for a list of the online utilities and the control file name associated with each utility.

INDSN(*data-set-name* **(***member-name***))**
Specifies what data set contains the utility statements and control statements. Do not specify a data set that contains double-byte character set data.

(*data-set-name***)**
Specifies the name of the data set.

If you do not specify a data set name, the default command procedure prompts you for the data set name.

(*member-name***)**
Specifies the member name.

You must specify the member name if the data set is partitioned.

CONTROL(*control-option***: ...)**
Specifies whether to trace the CLIST command execution.

NONE Omits tracing.

The **default** is **NONE**.

control-option
Lists one or more of the options given below. Separate items in the list by colons (:). To abbreviate, specify only the first letter of the option.

LIST Displays TSO commands after symbolic substitution and before command execution

CONLIST
Displays CLIST commands after symbolic substitution and before command execution

SYMLIST
Displays all executable statements (TSO commands and CLIST statements) before the scan for symbolic substitution

NONE Generates a CONTROL statement with the options NOLIST, NOCONLIST, and NOSYMLIST.

DB2I
Indicates the environment from which the DSNU CLIST command is called.

(NO) Indicates that DSNU CLIST command is not being called from the DB2I environment.

The **default** is **NO**.

(YES) Indicates that DSNU CLIST command is called from the DB2I environment and that card information can be used. Only the utility panels should execute the CLIST command with DB2I(YES).

DISCDSN(*data-set-name***)**
The cataloged data set name used by LOAD and REORG as a discard data set. For LOAD, this data set holds records not loaded; for REORG, it holds records not reloaded.

PUNCHDSN(*data-set-name***)**
The cataloged data set name used by REORG or UNLOAD to hold the generated LOAD utility control statements for UNLOAD EXTERNAL or DISCARD.

COPYDSN(*data-set-name***)**
The name of the cataloged data set that is used as a target (output) data set. If you do not supply this information, the CLIST command prompts you for it. It is optional for LOAD and for REORG with SHRLEVEL NONE; it is required for COPY, for MERGECOPY, and for REORG with SHRLEVEL REFERENCE or CHANGE.

COPYDSN2(*data-set-name***)**
The name of the cataloged data set that is used as a target (output) data set for the backup copy. It is optional for COPY, MERGECOPY, LOAD, and REORG.

RCPYDSN1(*data-set-name***)**
The name of the cataloged data set that is used as a target (output) data set for the remote site primary copy. It is optional for COPY, LOAD, and REORG.

RCPYDSN2(*data-set-name***)**
The name of the cataloged data set that is used as a target (output) data set for the remote site backup copy. It is optional for COPY, LOAD, and REORG.

RECDSN(*data-set-name***)**
The cataloged data set name that is used by LOAD for input or by REORG TABLESPACE or UNLOAD as the unload data set. If you do not supply this information, the CLIST command prompts you for it. It is required for LOAD and REORG TABLESPACE only.

EDIT
Specifies whether to invoke an editor to edit the temporary file generated by the CLIST command.

(NO)
Does not invoke an editor.

The **default** is **NO**.

(SPF)
Invokes the ISPF editor.

(TSO)
Invokes the TSO editor.

RESTART
Specifies whether this job restarts a current utility job, and, if so, at what point it is to be restarted.

(NO) Indicates that the utility is a new job, not a restarted job. The utility identifier (UID) must be unique for each utility job step.

The **default** is **NO**.

(CURRENT)
Restarts the utility at the last commit point.

(PHASE)
Restarts the utility at the beginning of the current stopped phase. You can determine the current stopped phase using the DISPLAY UTILITY command.

(PREVIEW)
Restarts the utility in preview mode.

SUBMIT
Specifies whether to submit the generated JCL for processing.

(NO) Does not submit the JCL data set for processing.

The **default** is **NO**.

(YES) Submits the JCL data set for background processing, using the TSO SUBMIT command.

(PROMPT)
Prompts you, after the data set is processed, to specify whether to submit the JCL data set for batch processing. You cannot use PROMPT when the CLIST command is executed in the TSO batch environment.

SYSTEM(*subsystem-name***)**
Specifies the DB2 subsystem or group attach name.

The **default** is **DSN**.

UID(*utility-id***)**
Provides a unique identifier for this utility job within DB2.

The **default** is *tso-userid.control-file-name*, where *control-file-name* for each of the utilities is listed below:

Utility	*control-file-name*
CHECK INDEX	DSNUCHI
CHECK DATA	DSNUCHD
CHECK LOB	DSNUCHL
COPY	DSNUCOP
DIAGNOSE	DSNUDIA
LOAD	DSNULOA
MERGECOPY	DSNUMER
MODIFY	DSNUMOD
QUIESCE	DSNUQUI
REBUILD INDEX	DSNUREB
RECOVER	DSNUREC
REORG INDEX	DSNURGI
REORG LOB	DSNURGL
REORG TABLESPACE	DSNURGT
REPAIR	DSNUREP
REPORT	DSNURPT
RUNSTATS	DSNURUN
STOSPACE	DSNUSTO
UNLOAD	DSNUUNL

UNIT(*unit-name***)**
Assigns a unit address, a generic device type, or a user-assigned group name for a device on which a new temporary or permanent data set resides. *unit-name* is placed after the UNIT clause of the generated DD statement.

The **default** is **SYSDA**.

VOLUME(*vol-ser*)
Assigns the serial number of the volume on which a new temporary or permanent data set resides. *vol-ser* is placed after the VOL=SER clause of the generated DD statement. If you omit VOLUME, the VOL=SER clause is omitted from the generated DD statement.

Reviewing DSNU CLIST command output

DSNU builds a one-step job stream. The JCL data set consists of a JOB statement, an EXEC statement that executes the DB2 utility processor and the required DD statements. This JOB statement also includes the SYSIN DD * job stream, as shown in Figure 4. Any of these statements can be edited.

```
//DSNUCOP JOB your-job-statement-parameters
//          USER=userid,PASSWORD=userword
//*ROUTE PRINT routing-information
//UTIL      EXEC  DSNUPROC,SYSTEM=DSN,UID=TEMP,UTPROC='
//SYSCOPY   DD    DSN=MYCOPIES.DSN8D71A.JAN1,DISP=(MOD,CATLG,CATLG),
//          UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSIN     DD    *
     COPY TABLESPACE DSN8D71A.DSN8S71D
          FULL NO
          SHRLEVEL CHANGE

/*
```

Figure 4. Control file DSNUCOP.CNTL. This is an example of the JCL data set before editing.

The following list describes the required JCL data set statements:

Statement
Description

JOB The CLIST command uses any JOB statements that you saved when using DB2I. If no JOB statements exist, DB2 produces a skeleton JOB statement that you can modify after the JCL is complete. The job name is DSNU, followed by the first three letters of the utility name you are using.

EXEC The CLIST command builds the EXEC statement. The values you specified for SYSTEM (DSN, by default), UID(TEMP), and RESTART (none) become the values of SYSTEM, UID, and UTPROC for the DSNUPROC.

The CLIST command builds the necessary JCL DD statements. Those statements vary depending on the utility that you execute. Data sets that might be required are listed under “Data sets used by online utilities” on page 24.

SYSPRINT DD SYSOUT=A
Utility messages are sent to the SYSPRINT data set. The generated JCL defines OUTPUT, SYSPRINT as SYSOUT=A. You can use the TSO command to control the disposition of the SYSPRINT data set. For example, you can send the data set to your terminal. For further information, see *OS/390 TSO/E Command Reference.*

UTPRINT DD SYSOUT=A
If any utility requires a sort, it executes DFSORT. Messages from that program are sent to UTPRINT. The generated JCL defines UTPRINT as SYSOUT=A.

SYSIN DD *
To build the SYSIN DD * job stream, DSNU copies the data set named by the

INDSN parameter. The INDSN data set does not change, and you can reuse it when the DSNU procedure has completed.

Editing the generated JCL data set

You can edit the data set before you process it by using the EDIT parameter on the command procedure. Use the editor to add a JCL statement to the job stream, change JCL parameters (such as *ddnames*), or change utility control statements.

If you use a *ddname* that is not the default on some utility statement that you use, you must change the *ddname* in the JCL generated by the DSNU procedure. For example, in the REORG TABLESPACE utility the default option for UNLDDN is SYSREC, and DSNU builds a SYSREC DD statement for REORG TABLESPACE. If you use a different value for UNLDDN, you must edit the JCL data set and change SYSREC to the *ddname* that you used.

When you finish editing the data set, you can either save changes to the data set (by issuing SAVE), or instruct the editor to ignore all changes.

The SUBMIT parameter specifies whether to submit the data set statement as a background job. The temporary data set that holds the JCL statement is reused. If you want to submit more than one job that executes the same utility, you must rename the JCL data sets and submit them separately.

Examples

Example 1: The following CLIST command statement generates a data set called *authorization-id*.DSNURGT.CNTL that contains JCL statements that invoke the DSNUPROC procedure.

```
%DSNU UTILITY(REORG TABLESPACE) INDSN(MYREOR.DATA)
  RECDSN(MYREOR.WORK) RESTART(NO)
  EDIT(TSO) SUBMIT(YES)
```

The DSNUPROC procedure invokes the REORG TABLESPACE utility. The MYREOR.DATA data set is merged into the JCL data set as SYSIN input. MYREOR.WORK is a temporary data set required by REORG TABLESPACE. The TSO editor is invoked to allow editing of the JCL data set, *authorization-id*.DSNURGT.CNTL. The TSO editor then submits the JCL data set as a batch job and will not be modified by this CLIST command statement until a new request is made to execute the REORG TABLESPACE utility.

Example 2: The following example shows how to invoke the CLIST command for the COPY utility.

```
%DSNU
  UTILITY (COPY)
  INDSN ('MYCOPY(STATEMNT)')
  COPYDSN ('MYCOPIES.DSN8D71A.JAN1')
  EDIT (TSO)
  SUBMIT (YES)
  UID (TEMP)
  RESTART (NO)
```

Using the supplied JCL procedure (DSNUPROC)

Another method of invoking a DB2 online utility uses the supplied JCL procedure, DSNUPROC, shown in Figure 5 on page 40. This procedure uses the parameters that you supply to build an appropriate EXEC statement that executes an online utility.

To execute the DSNUPROC procedure, you must write and submit a JCL data set like that built by the DSNU CLIST command, and shown in Figure 4 on page 36. In your JCL, the EXEC statement executes the DSNUPROC procedure. “DSNUPROC syntax” explains the parameters you can supply to that procedure and the syntax.

DSNUPROC syntax

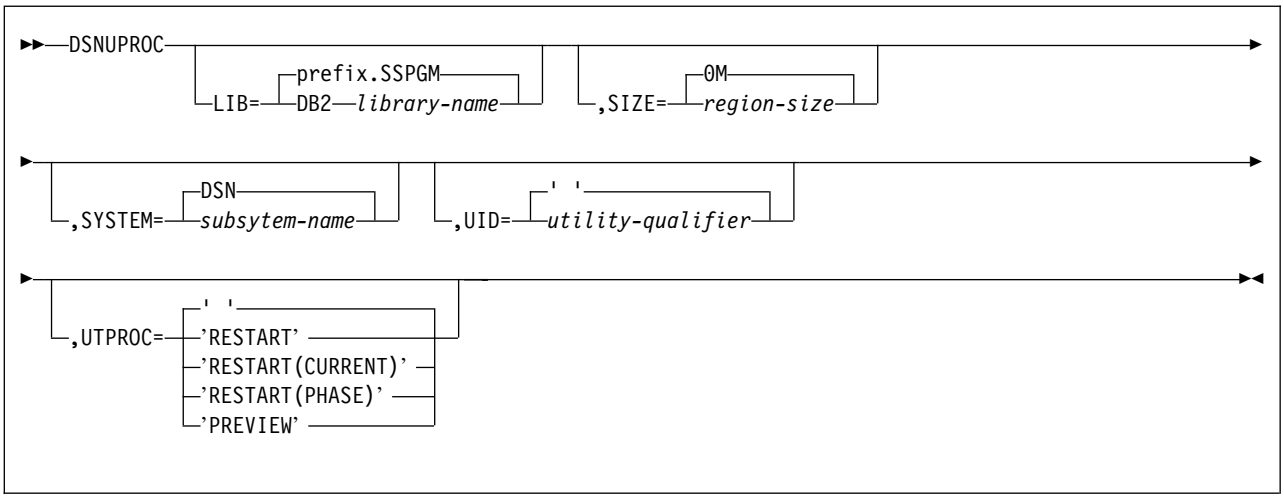

DSNUPROC option descriptions

The following list describes all the parameters. For example, in Figure 4 on page 36, you need to use only one parameter, UID=TEMP; for all others, you can use the defaults.

LIB= Specifies the data set name of the DB2 subsystem library.

The **default** is *prefix*.**SSPGM**.

SIZE=

Specifies the region size of the utility execution area; that is, the number of bytes of virtual storage allocated to this utility job.

The **default** is **0M**.

SYSTEM=

Specifies the DB2 subsystem or group attach name.

The **default** is **DSN**.

UID= Specifies the unique identifier for your utility job. The maximum name length is 16 characters. If the name contains special characters, enclose the entire name between apostrophes (for example, 'PETERS.JOB').

The **default** is an empty string.

UTPROC=

Controls restart processing.

The **default** is an empty string. Use the default if you do not want to restart a stopped job.

To restart the utility, specify:

'RESTART'
: To restart at the last commit point

'RESTART(CURRENT)'
: To restart at the last commit point.

'RESTART(PHASE)'
: To restart at the beginning of the phase executed last.

'PREVIEW'
: To restart in preview mode.

The procedure provides the SYSPRINT and UTPRINT DD statements for printed output. You must provide DD statements for SYSIN and other data sets that your job needs. See “Data sets used by online utilities” on page 24 for a description of data sets that you might need.

Figure 5 on page 40 is the DSNUPROC procedure that was executed by the JCL example in Figure 4 on page 36.

Sample DSNUPROC listing

```
//DSNUPROC PROC LIB='DSN.SSPGM',SIZE=250K,SYSTEM=DSN,UID='',UTPROC=''
//*
//*********************************************************************
//*                                                                   *
//* PROCEDURE-NAME:      DSNUPROC                                     *
//*                                                                   *
//* DESCRIPTIVE-NAME:    UTILITY PROCEDURE                            *
//*                                                                   *
//*   COPYRIGHT = 5740-XYR (C) COPYRIGHT IBM CORP 1982, 1985          *
//*   OCO SOURCE MATERIALS. IBM CONFIDENTIAL                          *
//*   (IBM CONFIDENTIAL RESTRICTED WHEN AGGREGATED)                   *
//*   REFER TO COPYRIGHT INSTRUCTIONS FORM NUMBER G120-2083           *
//*                                                                   *
//*   STATUS = RELEASE 2, LEVEL 0                                     *
//*                                                                   *
//* FUNCTION:  THIS PROCEDURE INVOKES THE DB2 UTILITIES IN THE        *
//*             BATCH ENVIRONMENT                                     *
//*                                                                   *
//* PROCEDURE-OWNER:     UTILITY COMPONENT                            *
//*                                                                   *
//* COMPONENT-INVOKED:   DB2 UTILITIES (ENTRY POINT DSNUTILB).        *
//*                                                                   *
//* ENVIRONMENT:         TSO BATCH                                    *
//*                                                                   *
//* INPUT:                                                            *
//*      PARAMETERS:                                                  *
//*           LIB    = THE DATA SET NAME OF THE DB2  PROGRAM LIBRARY.  *
//*                    THE DEFAULT LIBRARY NAME IS DSN.SSPGM.         *
//*           SIZE   = THE REGION SIZE OF THE UTILITIES EXECUTION AREA.*
//*                    THE DEFAULT REGION SIZE IS 250K.               *
//*           SYSTEM = THE SUBSYSTEM NAME USED TO IDENTIFY THIS JOB    *
//*                    TO DB2 .  THE DEFAULT IS "DSN".                *
//*           UID  = SPECIFIES A UTILITY IDENTIFER TO UNIQUELY IDENTIFY*
//*                    A NON-TERMINATED UTILITY JOB WITHIN DB2.  THIS  *
//*                    IDENTIFIER IS PASSED TO DB2 VIA THE PARM=       *
//*                    PARAMETER OF THE EXEC CARD.  THE DEFAULT IS TO  *
//*                    OMIT THE PARAMETER.  THIS WILL CAUSE THE UTILITY*
//*                    FUNCTION TO USE ITS DEFAULT WHICH IS: USERID.   *
//*                    JOBNAME.                                        *
```

Figure 5. Sample listing of supplied JCL procedure DSNUPROC (Part 1 of 2)

```
//*           UTPROC = AN OPTIONAL INDICATOR USED TO INVOKE       @L415*
//*                    THE PREVIEW FACILITY OR TO INDICATE THAT   @L415*
//*                    THE USER WISHES TO INITIALLY START THE REQUESTED*
//*                    UTILITY OR TO RESTART A PREVIOUS EXECUTION OF   *
//*                    THE UTILITY.  IF OMITTED, THE UTILITY WILL      *
//*                    BE INITIALLY STARTED.  VALID VALUES ARE:   @L415*
//*                       PREVIEW = INVOKE UTILITY PREVIEW        @L415*
//*                                                                     *
//*                       RESTART = RESTART THE UTILITY AT THE LAST    *
//*                                 OR CURRENT COMMIT POINT.            *
//*                                                                     *
//*                       RESTART(CURRENT) = SAME AS "RESTART"    @L415*
//*                                                                     *
//*                       RESTART(PHASE) = RESTART THE UTILITY AT @L415*
//*                                 THE BEGINNING OF THE LAST       @L415*
//*                                 SUCCESSFULLY COMPLETED PHASE OR @L415*
//*                                 THE BEGINNING OF THE JOB.       @L415*
//*                                 THE JOB.                             *
//*                                                                      *
//* OUTPUT: NONE.                                                        *
//*                                                                      *
//* EXTERNAL-REFERENCES: NONE.                                           *
//*                                                                      *
//* CHANGE-ACTIVITY:                                                     *
//*                                                                      *
//*   LI415 - RESTART(CURRENT) and PREVIEW                       @L415*
//*                                                                      *
//*********************************************************************
//*
//DSNUPROC EXEC PGM=DSNUTILB,REGION=&SIZE,
//         PARM=(&SYSTEM,&UID,&UTPROC)
//STEPLIB  DD   DSN=&LIB,DISP=SHR
//*
//*********************************************************************
//*                                                                    *
//*  THE FOLLOWING DEFINE THE UTILITIES' PRINT DATA SETS               *
//*                                                                    *
//*********************************************************************
//*
//SYSPRINT DD   SYSOUT=A
//UTPRINT  DD   SYSOUT=A
```

Figure 5. Sample listing of supplied JCL procedure DSNUPROC (Part 2 of 2)

Creating the JCL data set yourself

DB2 online utilities execute as standard OS/390 jobs. To execute the utility, you must supply the JOB statement required by your installation and the JOBLIB or STEPLIB DD statements required to access DB2. You must also include an EXEC statement and a set of DD statements. The EXEC statement is described in “Data sets used by online utilities” on page 24. For a description of the DD statements you might need, see “Data sets used by online utilities” on page 24.

We recommend using DSNUPROC to invoke a DB2 online utility, rather than creating the JCL yourself. For more information, see “Using the supplied JCL procedure (DSNUPROC)” on page 38.

EXEC statement

The EXEC statement can be a procedure that contains the required JCL, or it can be of the form:

```
//stepname  EXEC PGM=DSNUTILB,PARM='system,[uid],[utproc]'
```

where the brackets, [ ], indicate optional parameters. The parameters have the following meanings:

DSNUTILB
Specifies the utility control program. The program must reside in an APF-authorized library.

system
Specifies the DB2 subsystem.

uid The unique identifier for your utility job.

utproc The value of the UTPROC parameter in the DSNUPROC procedure. Specify this option only when you want to restart the utility job. Specify:

'RESTART'
To restart at the last commit point

'RESTART(CURRENT)'
To restart the utility at the last commit point.

'RESTART(PHASE)'
To restart at the beginning of the phase executed last.

'RESTART(PREVIEW)'
To restart the utility in preview mode.

For the example in Figure 5 on page 40, you write:

```
//stepname  EXEC PGM=DSNUTILB,PARM='DSN,TEMP'
```

Chapter 4. Monitoring and controlling online utilities

This section contains procedures and guidelines for monitoring utilities, running utilities concurrently, terminating utilities, and restarting utilities.

Monitoring utilities with the DISPLAY UTILITY command

The information under this heading, up to “Running utilities concurrently” on page 44 is General-use Programming Interface and Associated Guidance Information, as defined in “Notices” on page 725.

Use the DB2 DISPLAY UTILITY command to check the current status of online utilities. Figure 6 shows an example of the output generated by the DISPLAY UTILITY command. In the example output, DB2 returns a message that indicates the member name (**A**), utility name (**B**), identifier (**C**), phase (**D**), and status (**E**). The message indicates the number of pages or records that are processed by the utility (**F**) The output also reports the number of objects in the list (**G**) and the last object started (**H**).[1]The output might also report additional information for an executing utility, such as log phase estimates or utility subtask activity.

```
DSNU100I - DSNUGDIS - USERID = SAMPID
           A   MEMBER = DB1G
           C   UTILID = RUNTS
               PROCESSING UTILITY STATEMENT 1
           B   UTILITY = RUNSTATS
           D   PHASE = RUNSTATS   F  COUNT = 0
 G  NUMBER OF OBJECTS IN LIST = n
 H  LAST OBJECT STARTED = m
           E   STATUS = STOPPED
DSN9022I - DSNUGCC  '-DISPLAY UTILITY' NORMAL COMPLETION
```

Figure 6. DISPLAY UTILITY command sample output

Determining the status of a utility

An online utility can have one of these statuses:

Status (E)
Description

Active The utility has started execution.

Stopped
The utility has abnormally stopped executing before completion, but the table spaces and indexes accessed by the utility remain under utility control. To make the data available again, you must either:

- Correct the condition that stopped the utility, and restart the utility so that it runs to termination, or
- Terminate the utility with the DB2 TERM UTILITY command (see “Terminating an online utility with the TERM UTILITY command” on page 45).

1. In a data sharing environment, the number of records is current when the command is issued from the same member on which the utility is executing. When issued from a different member, the count may lag substantially.

© Copyright IBM Corp. 1983, 2001

Terminated
The utility has been requested to terminate by the DB2 TERM UTILITY command. If the utility has terminated, there is no message.

Determining which utility phase is currently executing

DB2 online utility execution is divided into phases. Each utility starts with the UTILINIT phase that performs initialization and set up. Each utility finishes with a UTILTERM phase that cleans up after processing has completed. The other phases of online utility execution differ, depending on the utility. See the "Execution Phases" section in the descriptions of each utility. Output from the DISPLAY UTILITY command shows the phase that is currently executing.

Determining why a utility failed to complete

If an online utility job completes normally, it issues return code 0. If it completes with warning messages, it issues return code 4. Return code 8 means that the job failed to complete. Return code 12 means that an authorization error occurred.

During execution of the utility, any of these problems can cause a failure:

- **Problem:** DB2 terminates the utility job step and any subsequent utility steps.

  Solution: Submit a new utility job to execute the terminated steps. Use the same utility identifier for the new job to ensure that no duplicate utility is running.
- **Problem:** DB2 does not execute the particular utility function, but prior utility functions are executed.

  Solution: Submit a new utility step to execute the function.
- **Problem:** DB2 places the utility function in the stopped state.

  Solution: Restart the utility job step at either the last commit point or the beginning of the phase by using the same utility identifier. Alternatively, terminate the job step (by using TERM UTILITY (uid)) and resubmit it.
- **Problem:** DB2 terminates the utility and issues return code 8.

  Solution: One or more objects may be in a restrictive or advisory status. See "Appendix C. Resetting an advisory or restrictive status" on page 677 for more information on resetting the status of an object.

  Alternatively, a DEADLINE condition in Online REORG might have terminated the reorganization.

Running utilities concurrently

Some online utilities allow other utilities and SQL statements to run concurrently on the same target object. The online utility descriptions in Section 2 of this book feature a section on compatibility and concurrency. The section on concurrency and compatibility includes the following information:

- For each target object on which the utility acts, the section outlines the claim classes that the utility must claim or drain. The section also outlines the restrictive state that the utility sets on the target object.
- For other online utilities, the section summarizes the compatibility of the utility with the same target object. If two actions are compatible on a target object, they can run simultaneously on that object in separate applications. If compatibility depends on particular options of a utility, that is also shown.

Utilities can use additional threads to support parallelism. Consider increasing the subsystem parameters that control concurrent threads. See *DB2 Installation Guide* for a description of MAX BATCH CONNECT and MAX USERS on installation panel DSNTIPE.

See Part 5 (Volume 2) of *DB2 Administration Guide* for a description of the claim classes and the use of claims and drains by online utilities.

Running online utilities in a data sharing environment

This section discusses considerations for running online utilities in a data sharing environment.

Submitting online utilities: When you submit a utility job, you must specify the name of the DB2 subsystem to which the utility is to attach or the group attach name. If you do not use the group attach name, the utility job must run on the MVS system where the specified DB2 subsystem is running. You must be sure that MVS runs the utility job on the appropriate MVS system. You must use one of several MVS installation-specific statements to make sure this happens. These include:

- For JES2 multi-access spool (MAS) systems, the following statement can be inserted into the utility JCL:

  ```
  /*JOBPARM SYSAFF=cccc
  ```

- For JES3 systems, the following statement can be inserted into the utility JCL:

  ```
  //*MAIN SYSTEM=(main-name)
  ```

Those JCL statements are described in *OS/390 MVS JCL Reference*. Your installation might have other mechanisms for controlling where batch jobs run, such as by using job classes.

Stopping and restarting utilities: In a data sharing environment, you can terminate an active utility (with the TERM UTILITY command) only on the DB2 subsystem on which it was started. If a DB2 subsystem fails while a utility is in progress, you must restart that DB2 and then you can terminate the utility from any system.

You can only restart a utility on a member that is running the same DB2 release level as the member on which the utility job was originally submitted. The same utility ID (UID) must be used to restart the utility. That UID is unique within a data sharing group. However, if DB2 fails, you must restart DB2 on either the same or another MVS system before you restart the utility.

Terminating an online utility with the TERM UTILITY command

The information under this heading, up to “Restarting an online utility” on page 46 is General-use Programming Interface and Associated Guidance Information, as defined in “Notices” on page 725.

Use the TERM UTILITY command to terminate the execution of an active utility or release the resources associated with a stopped utility.

After you issue the TERM UTILITY command, you cannot restart the terminated utility. It is possible that the objects on which the utility was operating might be left in an indeterminate state. In many cases you cannot rerun the same utility without first recovering the objects on which the utility was operating. The situation varies,

depending on the utility and the phase that was in process when you issued the command. These considerations are particularly important when terminating the COPY, LOAD, and REORG utilities.

If you cannot restart a utility, you might have to terminate it to make the data available to other applications. To terminate a utility, issue the DB2 TERM UTILITY command. Use the command only if you must start the utility from the beginning.

In a data sharing environment, TERM UTILITY is effective for active utilities when submitted from the DB2 subsystem that originally issued the command. You can terminate a stopped utility from any active member of the data sharing group.

If the utility is active, TERM UTILITY terminates it at the next commit point. It then performs any necessary cleanup operations.

You might choose to put TERM UTILITY in a conditionally executed job step; for example, if you never want to restart certain utility jobs. Figure 7 shows a sample job stream for our COPY example:

```
//TERM EXEC PGM=IKJEFT01,COND=((8,GT,S1),EVEN)
//*
//*************************************************************
//* IF THE PREVIOUS UTILITY STEP, S1, ABENDS, ISSUE A
//* TERMINATE COMMAND.  IT CAN NOT BE RESTARTED.
//*************************************************************
//*
//SYSPRINT DD  SYSOUT=A
//SYSTSPRT DD  SYSOUT=A
//SYSOUT   DD  SYSOUT=A
//SYSUDUMP DD  SYSOUT=A
//SYSTSIN  DD  *
DSN SYSTEM(DSN)
-TERM UTILITY(TEMP)
END
/*
```

Figure 7. Example of conditionally executed TERM UTILITY

Alternatively, consider specifying the TIMEOUT TERM parameter for some Online REORG situations.

Restarting an online utility

If a utility finishes abnormally, it is often possible to restart it. With the restart procedure, you avoid repeating much of the work that had already been done.

Two different methods of restart are available:

Phase restart can be done from the beginning of the phase that was being processed.

Current restart can be done from the last checkpoint taken during the execution of the utility phase. If the utility phase does not take any checkpoints or has not reached the first checkpoint, current restart is equivalent to phase restart.

Updating the JCL data set for restarting a utility

To restart a DB2 online utility, update the original JCL data set with the RESTART parameter. This can be accomplished using one of three methods:

- **Using DB2I**. Restart the utility following these steps:

1. Access the DB2 UTILITIES panel.
2. Fill in the panel fields as documented in Figure 2 on page 30, except for field 5.
3. Change field 5 to CURRENT or PHASE, depending on the method of restart desired.
4. Press ENTER.

- **Using the DSNU CLIST command**. Restart the utility by invoking the DSNU CLIST command as described in "Using the DSNU CLIST command in TSO" on page 31, but change the value of the RESTART parameter, using either RESTART, RESTART (CURRENT) or RESTART(PHASE).
- **Creating your own JCL**. If you create your own JCL you must specify RESTART (CURRENT) or RESTART(PHASE) to restart the utility. You must also check the DISP parameters on the DD statements. For example, DD statements that have DISP=NEW and need to be reused must have DISP changed to OLD or MOD. If generation data groups (GDGs) are used and any (+1) generations were cataloged, then ensure that the JCL is changed to GDG (+0) for such data sets.

  Automatically generated JCL normally has DISP=MOD. DISP=MOD allows a data set to be allocated during the first execution and then reused during a restart.

  Note that for templates on restart, DB2 automatically changes the disposition from NEW to MOD. Therefore, no changes are needed to template specifications for restart.

Adding or deleting utility statements

Restart processing remembers the relative position of the utility statement in the input stream. Therefore, you must include all the original utility statements when restarting any online utility; however, you can add or delete DIAGNOSE statements.

Restarting after the output data set is full

Special considerations exist when restarting a utility at the last commit point after the output data set runs out of space (for example, ABENDx37 on SYSUT1).

If you receive an out-of-space condition on the output data set, follow these steps before restarting the utility at the last commit point:

1. Copy the output data set to a temporary data set. Use the same DCB parameters. Use MVS utilities that do not reblock the data set during the copy operation (for example, DFDSS ADRDSUU or DFSORT ICEGENER). Avoid using the MVS utilities IEBGENER or ISPF 3.3.
2. Delete or rename the output data set. Redefine the data set with additional space. Use the same VOLSER (if the data set is not cataloged), the same DSNAME, and the same DCB parameters. You should know the current DCB parameters before attempting to redefine the data set.
3. Copy the data from the temporary data set into the new, larger output data set. Use MVS utilities that do not reblock the data set during the copy operation (for example, DFDSS ADRDSUU or DFSORT ICEGENER).

RESTARTING with TEMPLATES

Unlike other utility control statements, TEMPLATE control statements may be modified before RESTART, and in some cases they must be modified in order to correct a prior failure.

TEMPLATE allocation during a restart automatically adjusts data set dispositions to reallocate the data sets from the prior execution. No modification to the TEMPLATE DISP is required. If the prior failure was due to space problems on a data set, the same restart considerations apply as if DD cards were being used. If the prior failure was due to insufficient space on a volume, the TEMPLATE may be altered to point to a different volume or to reduce the size of the data set. DB2 checkpoints the values used for TEMPLATE DSN variables and the old values will be reused on restart.

RESTARTING with LISTS

Unlike other utility control statements, LISTDEF control statements may be modified before RESTART. However, their modification will not affect the currently running utility. It will only affect those utilities that follow it.

If the utility being restarted was processing a LIST you will see a list size greater than 1 on the DSNU100 or DSNU105 message. DB2 checkpoints the expanded, enumerated list contents prior to executing the utility. This checkpointed list will be used to restart the utility at the point of failure. After a successful restart, the LISTDEF is re-expanded before it is used by later utilities in the same job step.

Other restart hints

The following guidelines provide additional information for restarting utilities:

- The VOLSER (if the data set is not cataloged) and the data set name are used to track utility restart information. They must be the same as what you specified in the original JCL data set for the utility you want to restart. If you copy a work data set, such as SYSUT1, after an ABENDB37 and the number of volumes changed, do not specify RESTART CURRENT. It will result in an ABEND413–1C. To prevent this abend, start the utility in RESTART(PHASE).
- When restarting a utility with cataloged data sets, do not specify VOL=SER=. Let DB2 determine the VOLSER of the data sets from the system catalog.
- Do not change the utility function that is currently stopped and the DB2 objects on which it is operating. However, you can change other parameters that are related to the stopped step and subsequent utility steps.
- Do not specify MVS automatic step restart.
- If a utility is restarted in the UTILINIT phase, it is re-executed from the beginning of the phase.
- When you restart a LOAD, REBUILD INDEX, or REORG job in which you specified the STATISTICS keyword, DB2 does not collect inline statistics. You should run the RUNSTATS utility after the restarted utility completes. The exception is REORG UNLOAD PAUSE; When restarted after the pause, it collects statistics.
- If a utility abends or system failure occurs while the utility is in the UTILTERM phase, you must restart it with RESTART(PHASE).

Restart is not always possible. The restrictions applying to the phases of each utility are discussed under the description of each utility.

Chapter 5. CATMAINT

The CATMAINT utility updates the catalog; it should be run during migration to a new release of DB2 or when instructed to do so by the IBM service.

Output: Output for CATMAINT UPDATE is the updated catalog.

Authorization required: The authorization required for CATMAINT is installation SYSADM.

Execution phases of CATMAINT: The CATMAINT utility operates in these phases:

Phase Description
UTILINIT
Initialization and setup to the new type
UTILTERM
Cleanup

Syntax and options of the control statement

The utility control statement defines the function that the utility job performs. Use the ISPF/PDF edit function to create a control statement, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

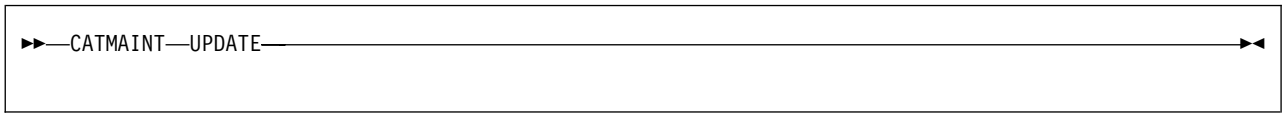

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

UPDATE
Indicates you want to update the catalog. Run this option only when you migrate to a new release of DB2.

Instructions for running CATMAINT

To run CATMAINT, you must:

1. Read “Before running CATMAINT” on page 50.
2. Prepare the necessary data sets, as described in “Data sets used by CATMAINT” on page 50.
3. Create JCL statements by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23.
4. Prepare a utility control statement that specifies the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 50. (For a complete description of the syntax and options for CATMAINT, see “Syntax and options of the control statement”.)

© Copyright IBM Corp. 1983, 2001

5. Check “Concurrency and compatibility” on page 51 if you want to run other jobs concurrently on the same target objects.
6. Plan for restarting CATMAINT if the job doesn't complete, as described in “Terminating or restarting CATMAINT”.
7. Run CATMAINT.

See “Chapter 3. Invoking DB2 online utilities” on page 23 for an explanation of methods you can use to execute DB2 utilities.

Before running CATMAINT

The work file database is used for CATMAINT sorting. Prior to executing the CATMAINT utility, calculate the size of the work file database.

To calculate the size of the work file database, see *DB2 Installation Guide*.

Data sets used by CATMAINT

Table 2 describes the data sets used by CATMAINT. Include DD statements in all data sets that are used by your job.

Table 2. Data sets used by CATMAINT

Data Set	Description	Required?
SYSIN	An input data set containing the utility control statement	Yes
SYSPRINT	An output data set for messages	Yes

Instructions for specific tasks

To perform the following task, specify the options and values for that task in your utility control statement.

Updating the catalog for a new release: When you migrate to a new release of DB2, you must update the catalog for the prior release to the new version. The DSNTIJTC job, described in *DB2 Installation Guide*, runs CATMAINT UPDATE to update the catalog. Run CATMAINT UPDATE only when you migrate to a new release. DB2 displays message DSNU777I at several points during CATMAINT execution.

If you have type 1 indexes on the catalog or directory, the CATMAINT UPDATE utility abnormally terminates. See “Migrating the DB2 Subsystem” in Part 2 of *DB2 Installation Guide* for the steps necessary to migrate to a new release.

If necessary, message DSNU776I or DSNU778I can give you information about why an abend occurred.

Terminating or restarting CATMAINT

You can terminate CATMAINT using the TERM UTILITY command, but it leaves the indexes in REBUILD-pending status. See “Resetting the REBUILD-pending status” on page 209 for information on resetting this status.

You cannot restart CATMAINT.

Concurrency and compatibility

Catalog and directory index availability: The catalog or directory indexes are not available while CATMAINT is running. This can cause other jobs to time out with message DSNT376I or message DSNT501I.

Chapter 6. CHECK DATA

The CHECK DATA online utility checks table spaces for violations of referential and table check constraints, and reports information about violations that are detected.
CHECK DATA also checks for consistency between a base table space and a LOB
table space.

Run CHECK DATA after a conditional restart or a point-in-time recovery on all table
spaces where parent and dependent tables might not be synchronized or where
base tables and auxiliary tables might not be synchronized. CHECK DATA can be run against a base table space only, not a LOB table space.

For a diagram of CHECK DATA syntax and a description of available options, see "Syntax and options of the control statement" on page 54. For detailed guidance on running this utility, see "Instructions for running CHECK DATA" on page 58.

Output: CHECK DATA optionally deletes rows that violate referential or table check constraints. A row that violates one or more constraints is copied, once, to an exception table.

If any violation of constraints is found, CHECK DATA puts the table space being checked in the CHECK-pending status.

On successful execution, CHECK DATA resets the CHECK-pending status.

Authorization required: To run this utility, the privilege set of this process must include one of the following:
- STATS privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority

An ID with installation SYSOPR authority can also execute CHECK DATA. However, you cannot use SYSOPR authority to execute CHECK DATA on table space SYSDBASE in database DSNDB06 or on any object except SYSUTILX in database DSNDB01.

If you specify the DELETE option, then the privilege set must include the DELETE privilege on the tables being checked. If you specify the FOR EXCEPTION option, then the privilege set must include the INSERT privilege on any exception table used. If you specify the AUXERROR INVALIDATE option, then the privilege set must include the UPDATE privilege on the base tables containing LOB columns.

Execution phases of CHECK DATA:

Phase	Description
UTILINIT	Initialization
SCANTAB	Extract foreign keys; use foreign key index if it matches exactly, else scan table
SORT	Sort foreign keys if not extracted from foreign key index
CHECKDAT	Look in primary indexes for foreign key parents, and issue messages to report errors detected
REPORTCK	Copy error rows into exception tables, and delete them from source table if DELETE YES is specified

© Copyright IBM Corp. 1983, 2001

UTILTERM Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

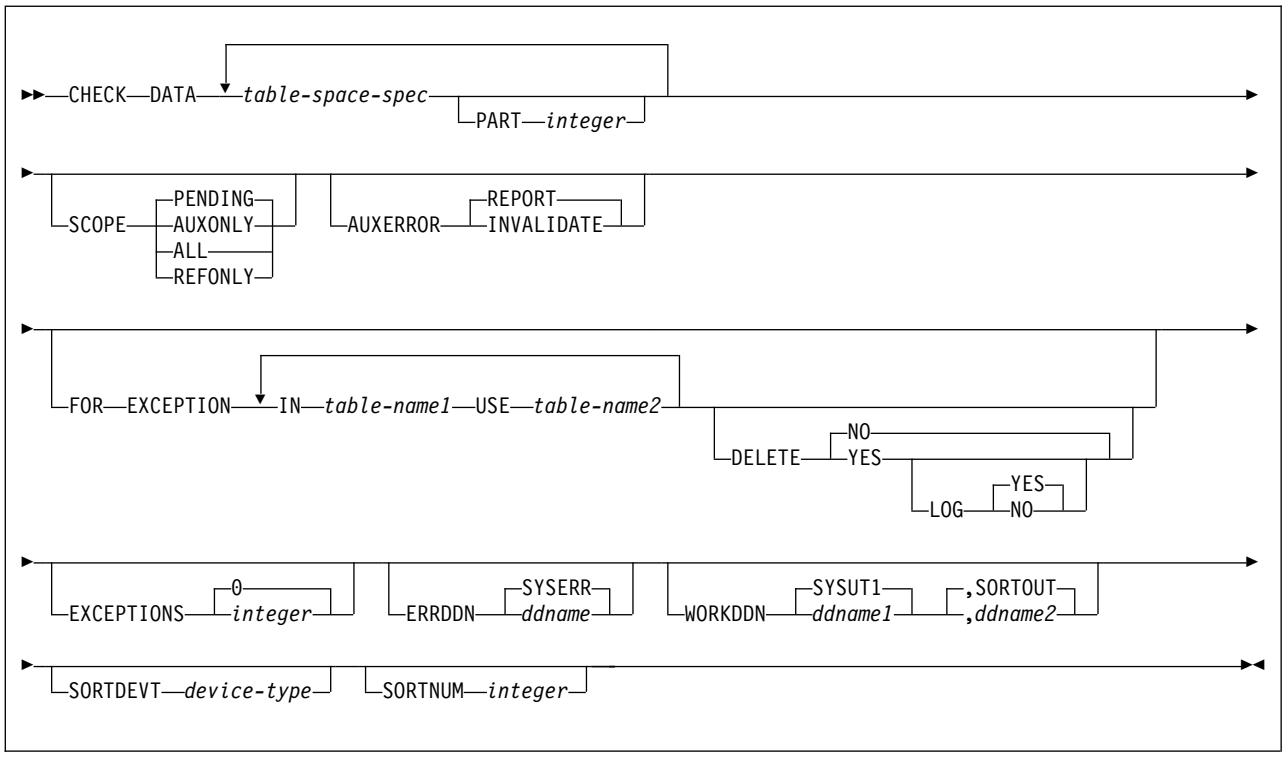

table-space-spec:

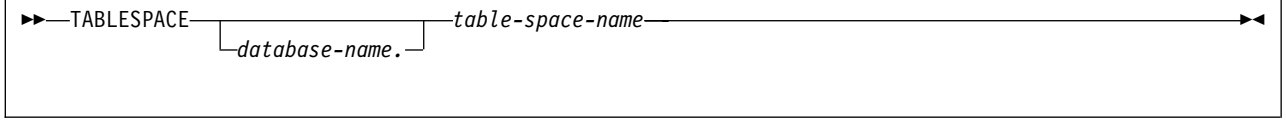

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

DATA Indicates that you are checking referential and table check constraints.

TABLESPACE *database-name.table-space-name*
Specifies the table space to which the data belongs.

database-name is the name of the database and is optional. The **default** is **DSNDB04**.

table-space-name is the name of the table space.

PART *integer* Identifies which partition to check for constraint violations.

integer is the number of the partition and must be in the range from 1 to the number of partitions defined for the table space. The maximum is 254.

SCOPE Limits the scope of the rows in the table space that are to be checked.

PENDING
Indicates that only those rows that are to be checked are those that are in table spaces, partitions, or tables that are in CHECK-pending status. The referential integrity check, constraint check, and the LOB check are all performed.

If you specify this option for a table space that is *not* in CHECK-pending status, the table space is ignored.

The **default** is **PENDING**.

AUXONLY
Indicates that only the LOB column check is to be performed for table spaces that have tables with LOB columns. The referential integrity and constraint checks are not performed.

ALL Indicates that all dependent tables in the specified table spaces are to be checked. The referential integrity check, constraint check, and the LOB check are performed.

REFONLY
Same as the **ALL** option, except the LOB column check is not performed.

AUXERROR Specifies the action that CHECK DATA is to perform when a LOB column check error is found.

REPORT A LOB column check error is reported with a warning message. The base table space is set to the auxiliary CHECK-pending (ACHKP) status.

The **default** is **REPORT**.

INVALIDATE A LOB column check error is reported with a warning message. The base table LOB column is set to an invalid status. The base table space is set to the auxiliary warning (AUXW) status. Before using CHECK DATA to check LOBs:

1. Run CHECK LOB to ensure the validity of the LOB table space.
2. Run REBUILD INDEX or CHECK INDEX on the index on the auxiliary table to ensure its validity.

CHECK DATA

FOR EXCEPTION

Indicates that any row in violation of referential or table check constraints is copied to an exception table.

If any row violates more than one constraint, it appears no more than once in the exception table.

IN *table-name1*

Specifies the table (in the table space specified on the TABLESPACE keyword) from which rows are to be copied.

table-name1 is the name of the table.

USE *table-name2*

Specifies the exception table into which error rows are to be copied.

table-name2 is the name of the exception table and must be a base table; it cannot be a view, synonym or alias.

DELETE Indicates whether or not rows in violation of referential or table check constraints are deleted from the table space. You can only use this option if you have used the FOR EXCEPTION keyword.

NO Indicates that error rows remain in the table space. Primary errors in dependent tables are copied to exception tables.

The **default** is **NO**.

If DELETE NO is specified, and constraint violations are detected, the table space is placed in the CHECK pending status.

YES Indicates that error rows are deleted from the table space. Deleted rows from both dependent and descendent tables are placed into exception tables.

LOG Specifies the logging action taken when records are deleted.

YES

Logs all records deleted during the REPORTCK PHASE.

NO

Does not log any records that are deleted during the REPORTCK phase. If any rows are deleted, CHECK DATA places the table space in the COPY-pending status, and places any indexes that were defined with the COPY YES attribute in the informational COPY-pending status.

Attention: Use the LOG NO option with caution. You cannot recover data across a point in the log in which CHECK DATA DELETE YES LOG NO was used.

EXCEPTIONS *integer*

Specifies maximum number of exceptions, which are reported by messages only. CHECK DATA terminates in the CHECKDAT phase when it reaches the number of exceptions specified; if termination occurs, the error rows are not written to the EXCEPTION table.

Only records containing *primary* referential integrity errors or table check constraint violations are applied toward the exception limit. There is no limit on the number of records containing secondary errors.

integer is the maximum number of exceptions. The **default** is **0**, which indicates no limit on the number of exceptions.

ERRDDN *ddname*

Specifies a DD statement for an error processing data set.

ddname is the DD name. The **default** is **ERRDDN SYSERR.**

The **ERRDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

WORKDDN (*ddname1***,***ddname2***)**

Specifies the DD statements for the temporary work file for sort input and the temporary work file for sort output. A temporary work file for sort input and output is *required.*

ddname1 is the DD name of the temporary work file for sort input.

The **default** is **SYSUT1**.

ddname2 is the DD name of the temporary work file for sort output.

The **default** is **SORTOUT**.

The **WORKDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, WORKDDN uses the DDNAME. For more information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

SORTDEVT *device-type*

Specifies the device type for temporary data sets to be dynamically allocated by DFSORT. It can be any device type acceptable to the DYNALLOC parameter of the SORT or OPTION control statement for DFSORT, as described in *DFSORT Application Programming: Guide.*

device-type is the device type. If you omit SORTDEVT and a sort is required, you must provide the DD statements that the sort program requires for the temporary data sets.

Do not use a TEMPLATE specification to dynamically allocate sort work data sets. The presence of the SORTDEVT keyword controls dynamic allocation of these data sets.

SORTNUM *integer*

Tells the number of temporary data sets to be dynamically allocated by the sort program.

integer is the number of temporary data sets.

If you omit SORTDEVT, SORTNUM is ignored.

If you use SORTDEVT and omit SORTNUM, no value is passed to DFSORT; it is allowed to take its own default.

Instructions for running CHECK DATA

To run CHECK DATA, you must:

1. Read “Before running CHECK DATA”.
2. Prepare the necessary data sets, as described in “Data sets used by CHECK DATA” on page 61.
3. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for CHECK DATA, see “Sample control statements” on page 66.)
4. Prepare a utility control statement that specifies the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 62. (For a complete description of the syntax and options for CHECK DATA, see “Syntax diagram” on page 54.)
5. Check the compatibility table in “Concurrency and compatibility” on page 65 if you want to run other jobs concurrently on the same target objects.
6. Plan for restarting CHECK DATA if the job doesn’t complete, as described in “Terminating or restarting CHECK DATA” on page 64.
7. Run CHECK DATA.

See “Chapter 3. Invoking DB2 online utilities” on page 23 for a description of ways to execute DB2 utilities.

Before running CHECK DATA

For a table with no LOB columns:

You should run CHECK INDEX on primary key indexes and foreign key indexes before running CHECK DATA to ensure that the indexes used by CHECK DATA are valid. This is especially important before using CHECK DATA with the DELETE YES or PART options.

For a table with LOB columns:

If you plan to run CHECK DATA on a base table space containing at least one LOB column, complete the following steps prior to running CHECK DATA:

1. Run CHECK LOB on the LOB table space.
2. Run CHECK INDEX on the index on the auxiliary table prior to running CHECK DATA to ensure the validity of the LOB table space and index on the auxiliary table.
3. Run CHECK INDEX on the base table space indexes.

If the LOB table space is in either the CHECK-pending or RECOVER-pending status, or if the index on the auxiliary table is in REBUILD-pending status, CHECK DATA will issue an error message and fail.

Create exception tables:

An *exception table* is a user-created table that duplicates the definition of a dependent table. The dependent table is checked with the CHECK DATA utility. It consists of at least *n* columns, where *n* is the number of columns of the dependent table. The CHECK DATA utility copies the deleted rows from the dependent table to the exception table. Table 3 describes the contents of an exception table.

Table 3. Exception tables

Column	Description	Required?	Data Type and Length	NULL Attribute
1 to n	Corresponds to columns in the table being checked. These columns hold data from rows in the table that violate referential or table check constraints.	Yes	The same as the corresponding columns in the table being checked.	The same as the corresponding columns in the table being checked.
n+1	Identifies the RIDs of the invalid rows of the table being checked.	No	CHAR(4); CHAR(5)[1] for table spaces defined with LARGE or DSSIZE options	Anything
n+2	Starting time of the CHECK utility	No	TIMESTAMP	Anything
≥ *n+2*	Additional columns that are not used by the CHECK utility	No	Anything	Anything

Notes:

1. You can use CHAR(5) for any type of table space.

If you delete rows using the CHECK DATA utility and use SCOPE ALL, you must create exception tables for all tables that are named in the table spaces and for all their descendents. All descendents of any row will be deleted.

When creating or using exception tables, be aware of the following:

- The exception tables should not have any unique indexes or referential or table check constraints that might cause errors when CHECK DATA inserts rows in them.
- You can create a new exception table before you run CHECK DATA, or use an existing exception table. The exception table can contain rows from multiple invocations of CHECK DATA.
- If column *n+2* is of type TIMESTAMP, CHECK DATA records the starting time. Otherwise, it does not use this column.
- You must have DELETE authorization on the dependent table being checked.
- You must have INSERT authorization on the exception table.
- Column names in the exception table can be given any name.
- Any change to the structure of the dependent table (such as a dropped column) is not automatically recorded in the exception table. You must make that change in the exception table.

Exception processing a table with a LOB column:

If you use exception tables, the exception table for the base table must have a similar LOB column and a LOB table space for each LOB column. If an exception is found, DB2 moves the base table row with its LOB column to the exception table, and moves the LOB column into the exception table's LOB table space. If you specify DELETE YES, DB2 deletes the base table row and the LOB column.

An auxiliary table cannot be an exception table. A LOB column check error is not included in the exception count. A row with a LOB column check error only does not participate in exception processing.

Example: creating an exception table for the project activity table

General-use Programming Interface

There is a clause of CREATE TABLE that makes the exception table easy to create. You can create an exception table for the project activity table by using these SQL statements:

```
CREATE TABLE EPROJACT
   LIKE DSN8710.PROJACT
   IN DATABASE DSN8D71A;

ALTER TABLE EPROJACT
   ADD RID CHAR(4);

ALTER TABLE EPROJACT
   ADD TIME TIMESTAMP NOT NULL WITH DEFAULT;
```

The first statement requires the SELECT privilege on table DSN8710.PROJACT and the privileges usually required to create a table.

Table EPROJACT has the same structure as table DSN8710.PROJACT, but it has two extra columns:

- Its first five columns mimic the columns of the project activity table; they have exactly the same names and descriptions. Although the column names are the same, they do not have to be. However, the rest of the column attributes for the initial columns must be same as those of the table being checked.
- The next column, added by ALTER TABLE, is optional; CHECK DATA uses it as an identifier. The name “RID” is an arbitrary choice—if the table already has a column with that name, you have to use a different name. The description of this column, CHAR(4), is required.
- The final timestamp column is also useful. If the timestamp column is defined, a row identifier (RID) column must precede this column. You might define a permanent exception table for each table that is subject to referential or table check constraints. You can define it once and use it to hold invalid rows detected by CHECK DATA. The TIME column allows you to identify rows that were added by the most recent run.

Eventually, you correct the data in the exception tables, perhaps with an SQL UPDATE, and transfer the corrections to the original tables by using statements such as:

```
INSERT INTO DSN8710.PROJACT
   SELECT PROJNO, ACTNO, ACSTAFF, ACSTDATE, ACENDATE
     FROM EPROJACT
     WHERE TIME > CURRENT TIMESTAMP - 1 DAY;
```

End of General-use Programming Interface

Complete all LOB column definitions

You must complete all LOB column definitions for a base table before running CHECK DATA. A LOB column definition is not complete until the LOB table space, auxiliary table and index on the auxiliary table have been created. If any LOB column definition is not complete, CHECK DATA will fail and issue error message DSNU075E.

Data sets used by CHECK DATA

Table 4 describes the data sets used by CHECK DATA. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 4. Data sets used by CHECK DATA

Data Set	Description	Required?
SYSIN	An input data set containing the utility control statement.	Yes
SYSPRINT	An output data set for messages.	Yes
Work data sets	Two temporary data sets for sort input and sort output. The symbolic names of the DD statement are specified with the WORKDDN option of the utility control statement. The default *ddname* for sort input is SYSUT1. The default *ddname* for sort output is SORTOUT. To find the approximate size in bytes of the work data sets, see page 61.	Yes
Error data set	An output data set that collects information about violations encountered during the CHECKDAT phase for referential constraints or the SCANTAB phase for check constraints. The symbolic name of the DD statement is specified with the ERRDDN parameter of the utility control statement. The default *ddname* is SYSERR.	Yes
UTPRINT	Contains messages from DFSORT (usually, SYSOUT or DUMMY).	No

The following objects are named in the utility control statement and do not require DD statements in the JCL:

Table space
Object to be checked. It is named in the CHECK DATA control statement and is accessed through the DB2 catalog. (If you want to check only one partition of a table space, use the PART option in the control statement.)

Exception tables
For each table in a table space that is checked, you can specify the name of an exception table in the utility control statement. Any row that violates a referential constraint is copied to the associated exception table. See page 58 for more information.

Defining work data sets: Three sequential data sets, described by the DD statements named in the WORKDDN and ERRDDN options, are required during execution of CHECK DATA.

To find the approximate size, in bytes, of the WORKDDN data set:
1. If a table space has a LOB column, count a total of 70 bytes for the LOB column, then go to step 3. If a table space does not have a LOB column, then go to step 2.
2. Add 17 to the length of the longest foreign key.
3. Multiply the sum by the number of keys and LOB columns checked.

Create the ERRDDN data set so that it is large enough to accommodate one error entry (length=60 bytes) per defect detected by CHECK DATA.

Creating the control statement

See “Syntax diagram” on page 54 for CHECK DATA syntax and option descriptions.
See “Sample control statements” on page 66 for examples of CHECK DATA usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

“Specify the scope of CHECK DATA”
“Checking several table spaces”
“Finding violations”
“Detecting and correcting constraint violations”
“Resetting CHECK-pending status” on page 63
“Interpreting LOB column errors” on page 63
“Resetting auxiliary CHECK-pending status” on page 64

Specify the scope of CHECK DATA

To specify the scope of CHECK DATA, it is normally sufficient to run CHECK DATA with SCOPE PENDING. DB2 keeps track of the data rows that must be checked to ensure the referential integrity of the table space. You should run SCOPE ALL whenever the scope information is in doubt. The scope information is recorded in the DB2 catalog. The scope information can become in doubt whenever you start the target table space with ACCESS(FORCE), or the catalog is recovered to a point in time.

If you only want to check tables with LOB columns, specify the AUXONLY option. If you want to check all dependent tables in the specified table spaces *except* tables with LOB columns, specify the REFONLY option.

Checking several table spaces

To check several table spaces, you can specify more than one table space in a CHECK DATA control statement. This technique is useful for checking a complete set of referentially related table spaces.

Finding violations

CHECK DATA issues a message for every row containing a referential or table check constraint violation. The violation is identified by:

- The RID of the row
- The name of the table that contained the row
- The name of the constraint being violated

Detecting and correcting constraint violations

To avoid problems, you should run CHECK DATA with DELETE NO to detect the violations before you attempt to correct the errors. If required, use DELETE YES after you analyze the output and understand the errors.

You can automatically delete rows that violate referential or table check constraints by specifying CHECK DATA with DELETE YES. However, you should be aware of the following possible problems:

- The violation might be created by a non-residential integrity error. For example, the indexes on a table might be inconsistent with the data in a table.

- Deleting a row might cause a cascade of secondary deletes in dependent tables. The cascade of deletes might be especially inconvenient within referential integrity cycles.
- The error might be in the parent table.
- Deleting a row might make the time error harder to detect.
- Valid rows might be detected.

CHECK DATA uses the primary key index and all indexes that match a foreign key exactly. Therefore, before running CHECK DATA, ensure that the indexes are consistent with the data by using CHECK INDEX.

Resetting CHECK-pending status

If you run CHECK DATA with the DELETE NO option and referential or table check constraint violations are found, the table space or partition is placed in CHECK-pending status.

Take one of the following actions to remove the CHECK-pending status:

- Use the DELETE NO option if no tables contain rows that violate referential or table check constraints.
- Use the DELETE YES option to remove all rows in violation of referential or table check constraints.

Interpreting LOB column errors

If you run CHECK DATA AUXERROR REPORT or INVALIDATE on a base table space containing at least one LOB column, the following errors might be reported:

Orphan LOBs: An orphan LOB column is a LOB found in the LOB table space but not referenced by the base table space. An orphan can result if you recover the base table space to a point in time prior to the insertion of the base table row or prior to the definition of the LOB column. An orphan can also result if you recover the LOB table space to a point in time prior to the deletion of a base table row.

Missing LOBs: A missing LOB column is a LOB referenced by the base table space, but the LOB is not in the LOB table space. A missing LOB can result if you recover the LOB table space to a point in time when the LOB column is not in the LOB table space. This could be a point in time prior to the first insertion of the LOB into the base table, or when the LOB column is null or has a zero length.

Out-of-synch LOBs: An out-of-synch LOB error occurs when DB2 detects a LOB that is found in both the base table and the LOB table space, but the LOB in the LOB table space is at a different level. An LOB column is also out-of-synch if the base table LOB column is null or has a zero length, but the LOB is found in the LOB table space. An out-of-synch LOB can occur anytime you recover the LOB table space or the base table space to a prior point in time.

Invalid LOBs: An invalid LOB is an uncorrected LOB column error found by a previous execution of CHECK DATA AUXERROR INVALIDATE.

Detecting LOB column errors: If you specify either CHECK DATA AUXERROR REPORT or AUXERROR INVALIDATE and a LOB column check error is detected, DB2 reports a message identifying the table, row, column, and type of error. Any additional actions depend on the option you specify for the AUXERROR parameter.

Actions performed with AUXERROR REPORT: DB2 sets the base table space to the auxiliary CHECK-pending (ACHKP) status. If CHECK DATA encounters only invalid LOB columns and no other LOB column errors, the base table space is set to the auxiliary warning status.

Actions performed with AUXERROR INVALIDATE: DB2 sets the base table LOB column to an invalid status, and sets the base table space to the auxiliary warning (AUXW) status. You can use SQL to update a LOB column in the AUXW status, however, any other attempt to access the column will result in a -904 SQL return code.

See “Appendix C. Resetting an advisory or restrictive status” on page 677 for information about the resetting the restrictive table space status.

Resetting auxiliary CHECK-pending status

If there are tables with LOB columns in a table space that is recovered to a point in time, then RECOVER TABLESPACE sets the auxiliary check-pending (ACHKP) status on the table space.

Use one of the following actions to remove the auxiliary CHECK-pending status if no inconsistencies are found:

- Use the SCOPE(ALL) option to check all dependent tables in the specified table space. The checks include: referential integrity constraints, table check constraints, and the existence of LOB columns.
- Use the SCOPE(PENDING) option to check tables spaces or partitions with CHKP status. The checks include: referential integrity constraints, table check constraints, and the existence of LOB columns.
- Use the SCOPE(AUXONLY) option to check for LOB columns.

If you specified the AUXERROR(INVALIDATE) option and inconsistencies are found, the AUXW status is issued. See “Appendix C. Resetting an advisory or restrictive status” on page 677 for information about the resetting the restrictive table space status.

Terminating or restarting CHECK DATA

Terminating CHECK DATA

When you terminate CHECK DATA, table spaces remain in the CHECK-pending status as they were when the utility was terminated. The CHECKDAT phase places the table space in the CHECK-pending status when an error is detected; at the end of the phase, the CHECK-pending status is reset if no errors were detected. The REPORTCK phase resets the CHECK-pending status if you specify the DELETE YES option.

For instructions on terminating an online utility, see “Terminating an online utility with the TERM UTILITY command” on page 45.

Restarting CHECK DATA

You can restart a CHECK DATA utility job, but it starts from the beginning again.

For instructions on restarting a utility job, see “Restarting an online utility” on page 46.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Claims and drains: Table 5 shows the claim classes in which CHECK DATA claims and drains and any restrictive status the utility sets on the target object.

Table 5. Claim classes of CHECK DATA operations. Use of claims and drains; restrictive state set on the target object.

TARGET OBJECTS	CHECK DATA DELETE NO	CHECK DATA DELETE YES	CHECK DATA PART DELETE NO	CHECK DATA PART DELETE YES
Table space or Partition	DW/UTRO	DA/UTUT	DW/UTRO	DA/UTUT
Partitioning Index or Index Partition	DW/UTRO	DA/UTUT	DW/UTRO	DA/UTUT
Nonpartitioning index	DW/UTRO	DA/UTUT		DR
Logical partition of index			DW/UTRO	DA/UTUT
Primary index	DW/UTRO	DW/UTRO	DW/UTRO	DW/UTRO
RI dependent and descendent table spaces and indexes		DA/UTUT		DA/UTUT
RI exception table spaces and indexes (FOR EXCEPTION only)	DA/UTUT	DA/UTUT	DA/UTUT	DA/UTUT

Legend:
- DA: Drain all claim classes, no concurrent SQL access
- DR: Drain the repeatable read class, no concurrent access for SQL repeatable readers
- DW: Drain the write claim class, concurrent access for SQL readers
- UTUT: Utility restrictive state, exclusive control
- UTRO: Utility restrictive state, read only access allowed
- Blank: Object is not affected by this utility

Table 6. Claim classes on a LOB table space and index on the auxiliary table for CHECK DATA operations. Use of claims and drains; restrictive states set on the target object.

TARGET OBJECTS	CHECK DATA DELETE NO	CHECK DATA DELETE YES
LOB table space	DW/UTRO	DA/UTUT
Index on the auxiliary table	DW/UTRO	DA/UTUT

Legend:
- DW: Drain the write claim class, concurrent access for SQL readers
- DA: Drain all claim classes, no concurrent SQL access
- UTRO: Utility restrictive state, read only access allowed
- UTUT: Utility restrictive state, exclusive control

When you specify CHECK DATA AUXERROR INVALIDATE, a drain-all is performed on the base table space, and the base table space is set UTUT.

Compatibility: The following utilities are compatible with CHECK DATA and can run concurrently on the same target object:

- DIAGNOSE
- MERGECOPY
- MODIFY
- REPORT
- STOSPACE
- UNLOAD (when CHECK DATA DELETE NO)

SQL operations and other online utilities are incompatible.

To run on DSNDB01.SYSUTILX, CHECK DATA must be the only utility in the job step and the only utility running in the DB2 subsystem.

The index on the auxiliary table for each LOB column inherits the same compatibility and concurrency attributes of a primary index.

Sample control statements

Example 1: CHECK DATA with DELETE. The following shows CHECK DATA JCL for checking and deleting.

```
//STEP1    EXEC DSNUPROC,UID='IUIQU1UQ.CHK1',
//         UTPROC='',
//         SYSTEM='V61A'
//SYSUT1   DD DSN=IUIQU1UQ.CHK3.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(8000,(200,20),,,ROUND)
//SYSERR   DD DSN=IUIQU1UQ.CHK3.SYSERR,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SORTOUT  DD DSN=IUIQU1UQ.CHK3.STEP1.SORTOUT,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SYSIN    DD *
CHECK DATA TABLESPACE DSN8D71A.DSN8S71D
           TABLESPACE DSN8D71A.DSN8S71E
     FOR EXCEPTION IN DSN8710.DEPT        USE DSN8710.EDEPT
                   IN DSN8710.EMP         USE DSN8710.EEMP
                   IN DSN8710.PROJ        USE DSN8710.EPROJ
                   IN DSN8710.PROJACT     USE DSN8710.EPROJACT
                   IN DSN8710.EMPPROJACT  USE DSN8710.EEPA
     DELETE YES
//*
```

Example 2: Control statement for deleting error rows. Check for and delete all constraint violations in table spaces DSN8D71A.DSN8S71D and DSN8D71A.DSN8S71E.

```
CHECK DATA TABLESPACE DSN8D71A.DSN8S71D
 TABLESPACE DSN8D71A.DSN8S71E
 FOR EXCEPTION IN DSN8710.DEPT USE DSN8710.EDEPT
               IN DSN8710.EMP USE DSN8710.EEMP
               IN DSN8710.PROJ USE DSN8710.EPROJ
               IN DSN8710.PROJECT USE DSN8710.EPROJECT
               IN DSN8710.EMPPROJECT USE DSN8710.EEMPPROJECT
 DELETE YES
```

Chapter 7. CHECK INDEX

The CHECK INDEX online utility tests whether indexes are consistent with the data they index, and issues warning messages when an inconsistency is found.

CHECK INDEX should be executed after a conditional restart or a point-in-time recovery on all table spaces whose indexes may not be consistent with the data.

It should also be used before CHECK DATA to ensure that the indexes used by CHECK DATA are valid. This is especially important before using CHECK DATA with DELETE YES. When checking an auxiliary table index, CHECK INDEX verifies that each LOB is represented by an index entry, and that an index entry exists for every LOB.

For a diagram of CHECK INDEX syntax and a description of available options, see "Syntax and options of the control statement". For detailed guidance on running this utility, see "Instructions for running CHECK INDEX" on page 69.

Output: CHECK INDEX generates several messages that show whether the indexes are consistent with the data. See Part 2 of *DB2 Messages and Codes* for more information about these messages.

For unique indexes, any two null values are taken to be equal, unless the index was created with the UNIQUE WHERE NOT NULL clause. In that case, if the key is a single column, it can contain any number of null values, and CHECK INDEX does not issue an error message.

CHECK INDEX issues an error message if there are two or more null values and the unique index was not created with the UNIQUE WHERE NOT NULL clause.

Authorization required: To execute this utility, the privilege set of this process must include one of the following:
- STATS privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority.

An ID with installation SYSOPR authority can also execute CHECK INDEX, but only on a table space in the DSNDB01 or DSNDB06 databases.

Execution Phases of CHECK INDEX:

Phase	Description
UTILINIT	Initialization and setup
UNLOAD	Unloading of index entries
SORT	Sorting of unloaded index entries
CHECKIDX	Scanning of data to validate index entries
UTILTERM	Cleanup.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

© Copyright IBM Corp. 1983, 2001

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

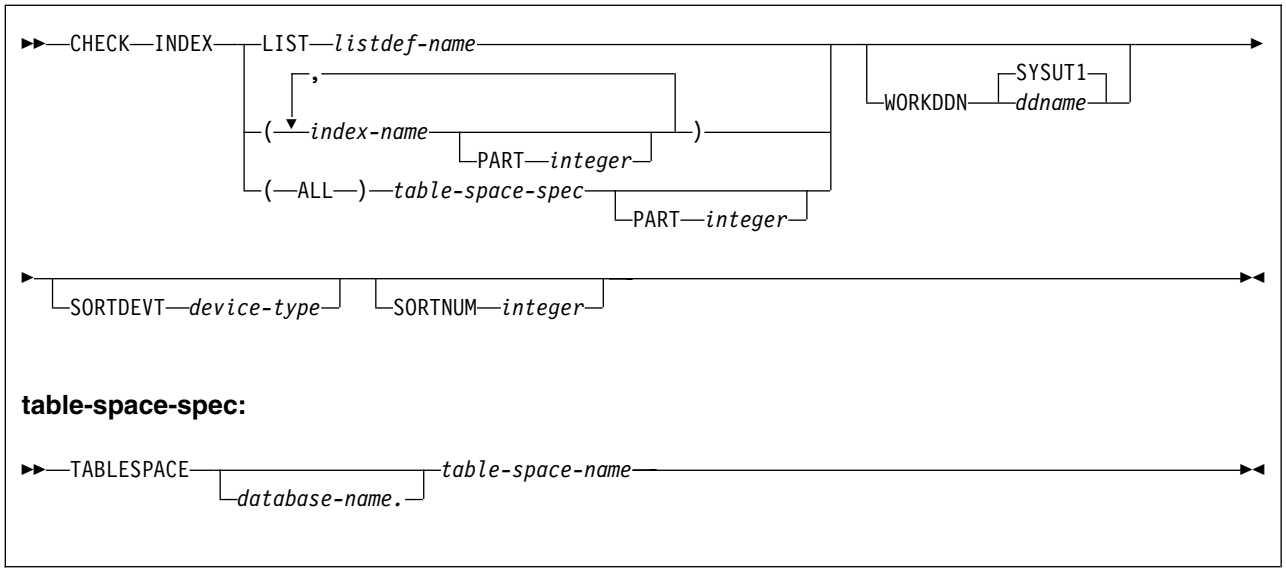

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

INDEX Indicates that you are checking for index consistency.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name. The list should contain only index spaces. The utility list manager will group indexes by their related table space and execute check index once per table space. CHECK INDEX allows one LIST keyword for each control statement in CHECK INDEX. Do not specify LIST with an *index-name* or a **TABLESPACE** specification. For more information about LISTDEF specifications, see “Generating object lists using LISTDEF” on page 144.

(*index-name***, ...)**
Specifies the indexes that are to be checked. All indexes must belong to tables in the same table space. If you omit this option, you must use the (ALL) TABLESPACE option. Then CHECK INDEX checks all indexes on all tables in the table space you specify.

index-name is the name of an index, in the form *creator-id.name*. If you omit the qualifier *creator-id.*, the user identifier for the utility job is used. If you use a list of names, separate items in the list by commas. Parentheses are required around a name or list of names.

PART *integer* Identifies a physical partition of a partitioning index or a logical partition of a nonpartitioning index to check.

integer is the number of the partition and must be in the range from 1 to the number of partitions defined for the table space. The maximum is 254.

(ALL) Specifies that all indexes in the specified table space referenced by the table space are to be checked.

TABLESPACE *database-name.table-space-name*
Specifies the table space from which all indexes will be checked. If an explicit list of index names is not given, then all indexes on all tables in the specified table space will be checked.

Do not specify TABLESPACE with an explicit list of index names.

database-name is the name of the database that the table space belongs to. The **default** is **DSNDB04**.

table-space-name is the name of the table space from which all indexes will be checked.

WORKDDN *ddname*
Specifies a DD statement for a temporary work file.

ddname is the DD name. The **default** is **SYSUT1**.

The **WORKDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

SORTDEVT *device-type*
Specifies the device type for temporary data sets to be dynamically allocated by DFSORT. It can be any device type acceptable to the DYNALLOC parameter of the SORT or OPTION control statement for DFSORT.

device-type is the device type. If you omit SORTDEVT and a sort is required, you must provide the DD statements that the sort program requires for the temporary data sets.

A TEMPLATE specification will not dynamically allocate sort work data sets. The SORTDEVT keyword controls dynamic allocation of these data sets.

SORTNUM *integer*
Tells the number of temporary data sets to be dynamically allocated by the sort program.

integer is the number of temporary data sets.

If you omit SORTDEVT, SORTNUM is ignored.

If you use SORTDEVT and omit SORTNUM, no value is passed to DFSORT; it is allowed to take its own default.

Instructions for running CHECK INDEX

To run CHECK INDEX, you must:

1. Prepare the necessary data sets, as described in 67.
2. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for CHECK INDEX, see “Sample control statements” on page 73.)

3. Prepare a utility control statement, specifying the options for the tasks you want to perform. For a complete description of the syntax and options for CHECK INDEX, see “Syntax and options of the control statement” on page 67.
4. Check the compatibility table in “Concurrency and compatibility” on page 72 if you want to run other jobs concurrently on the same target objects.
5. Plan for restart if the CHECK INDEX job doesn't complete, as described in “Terminating or restarting CHECK INDEX” on page 72.
6. Run CHECK INDEX.

See “Chapter 3. Invoking DB2 online utilities” on page 23 for an explanation of ways to execute DB2 utilities.

Data sets used by CHECK INDEX

Table 7 describes the data sets used by CHECK INDEX. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 7. Data sets used by CHECK INDEX

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
Work data set	A temporary data set for collecting index key values to be checked. Its DD name is specified with the WORKDDN option of the utility control statement. The default DD name is SYSUT1. To find the approximate size in bytes of the work data sets, see page 70.	Yes
UTPRINT	Contains messages from DFSORT (usually, SYSOUT or DUMMY).	No

The following object is named in the utility control statement and does not require a DD card in the JCL:

Index space
Object to be checked. It is named in the CHECK INDEX control statement and is accessed through the DB2 catalog. (If you want to check only one partition of an index, you must use the PART option in the control statement.)

Defining the work data set for CHECK INDEX: A single sequential data set, described by the DD statement specified in the WORKDDN option, is required during execution of CHECK INDEX.

To find the approximate size of the WORKDDN data set, in bytes:

1. For each table, multiply the number of records in the table by the number of indexes needing to be checked on the table.
2. Add the products obtained in step 1.
3. Add 9 to the length of the longest key.
4. Multiply the sum from step 2 by the sum from step 3.

Another method of estimating the size of the WORKDDN data set is to obtain the high-used relative byte address (RBA) for each index from a VSAM catalog listing. Then sum the RBAs.

Creating the control statement

See “Syntax diagram” on page 68 for CHECK INDEX syntax and option descriptions. See “Sample control statements” on page 73 for examples of CHECK INDEX usage.

Instructions for specific tasks

To perform the following task, specify the options and values documented with your utility control statement.

Checking a single logical partition

You can run CHECK INDEX on a single logical partition of a nonpartitioning index. However, there are some limitations on what CHECK INDEX can detect:

- It does not detect duplicate unique keys in different logical partitions. For example, logical partition 1 might have the following keys:

  ```
  A  B  E  F  T  Z
  ```

  and logical partition 2 might have these keys:

  ```
  M  N  Q  T  V  X
  ```

  In this example, the keys are unique within each logical partition, but both logical partitions contain the key, T; so for the index as a whole, the keys are not unique.
- It does not detect keys that are out of sequence between different logical partitions. For example, the following keys are out of sequence:

  ```
  1  7  5  8  9  10  12
  ```

  If keys 1, 5, 9 and 12 belong to logical partition 1 and keys 7, 8, and 10 belong to logical partition 2, then the keys within each partition are in sequence, but the keys for the index, as a whole, are out of sequence:

  ```
  LP 1   1     5     9      12
  LP 2      7     8     10
  ```

  When checking a single logical partition, this out of sequence condition is not detected.

Reviewing CHECK INDEX output

CHECK INDEX indicates whether or not a table space and its indexes are inconsistent, but does not correct any such inconsistencies. If CHECK INDEX detects inconsistencies, you should analyze the output to determine the problem and then correct the inconsistency. Perform the following actions to identify the inconsistency:

1. Examine the error messages from CHECK INDEX.
2. Verify the point in time (TOLOGPOINT, TORBA, or TOCOPY) for each object recovered. Use output from REPORT RECOVERY to determine a consistent point for both the table space and its indexes.
3. If the table space is correct, run the REBUILD INDEX utility to rebuild the indexes.
4. If the index is correct, determine a consistent point in time for the table space, and run the RECOVER utility on the table space. Run CHECK INDEX again to verify consistency.

5. If neither the table space nor its indexes are correct, determine a consistent point in time, then run the RECOVER utility job again, including the table space and its indexes all in the same list.

Terminating or restarting CHECK INDEX

CHECK INDEX can be terminated in any phase without any integrity exposure.

You can restart a CHECK INDEX utility job, although it starts over again from the beginning.

For instructions on restarting a utility job, see “Restarting an online utility” on page 46.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Claims and drains: Table 8 shows which claim classes CHECK INDEX claims and drains and any restrictive state the utility sets on the target object.

Table 8. Claim classes of CHECK INDEX operations. Use of claims and drains; restrictive states set on the target object.

Target	CHECK INDEX	CHECK INDEX PART
Table space or partition	DW/UTRO	DW/UTRO
Partitioning index or index partition	DW/UTRO	DW/UTRO
Nonpartitioning index	DW/UTRO	
Logical partition of an index		DW/UTRO

Legend:
- DW: Drain the write claim class, concurrent access for SQL readers
- UTRO: Utility restrictive state, read only access allowed
- Blank: Object is not affected by this utility

CHECK INDEX does not set a utility restrictive state if the target object is DSNDB01.SYSUTILX.

Compatibility: Table 9 shows which utilities can run concurrently with CHECK INDEX on the same target object. The target object can be a table space, an index space, or an index partition. If compatibility depends on particular options of a utility, that is also documented.

Table 9. CHECK INDEX compatibility

Action	CHECK INDEX
CHECK DATA	No
CHECK INDEX	Yes
CHECK LOB	Yes
COPY INDEXSPACE	Yes

Table 9. CHECK INDEX compatibility (continued)

Action	CHECK INDEX
COPY TABLESPACE	Yes
DIAGNOSE	Yes
LOAD	No
MERGECOPY	Yes
MODIFY	Yes
QUIESCE	Yes
REBUILD INDEX	No
RECOVER INDEX	No
RECOVER TABLESPACE	No
REORG INDEX	No
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No
REORG TABLESPACE UNLOAD ONLY or EXTERNAL	Yes
REPAIR DUMP or VERIFY	Yes
REPAIR DELETE or REPLACE	No
REPORT	Yes
RUNSTATS	Yes
STOSPACE	Yes
UNLOAD	Yes

To run on SYSIBM.DSNLUX01 or SYSIBM.DSNLUX02, CHECK INDEX must be the only utility within the same job step.

Sample control statements

Example 1: Check all indexes in a sample table space.

```
//STEP1   EXEC DSNUPROC,UID='IUIQU1UQ.CHK1',
//         UTPROC='',
//         SYSTEM='V71A'
//SYSUT1   DD DSN=IUIQU1UQ.CHK3.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(8000,(200,20),,,ROUND)
//SYSERR   DD DSN=IUIQU1UQ.CHK3.SYSERR,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SORTOUT  DD DSN=IUIQU1UQ.CHK3.STEP1.SORTOUT,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SYSIN    DD *
CHECK INDEX (ALL) TABLESPACE DSN8D71A.DSN8S71E
//*
```

Example 2: Check one index. Check the project-number index (DSN8710.XPROJ1) on the sample project table.

```
CHECK INDEX (DSN8710.XPROJ1)
  SORTDEVT SYSDA
```

Example 3: Check more than one index. Check the indexes DSN8710.XEMPRAC1 and DSN8710.XEMPRAC2 on the employee to project activity sample table.

```
CHECK INDEX NAME (DSN8710.XEMPRAC1, DSN8710.XEMPRAC2)
```

Example 4: Check all indexes on a table space. Check all indexes on the employee-table table space (DSN8S71E).

```
CHECK INDEX (ALL) TABLESPACE DSN8S71E
  SORTDEVT 3380
```

Chapter 8. CHECK LOB

The CHECK LOB online utility can be run against a LOB table space to identify any structural defects in the LOB table space and any invalid LOB values.

Run the CHECK LOB online utility against a LOB table space that is marked CHECK pending (CHKP) to identify structural defects. If none is found, the CHECK LOB utility turns the CHKP status off.

Run the CHECK LOB online utility against a LOB table space that is in auxiliary warning (AUXW) status to identify invalid LOBs. If none exists, the CHECK LOB utility turns AUXW status off.

Run CHECK LOB after a conditional restart or a point-in-time recovery on all table spaces where LOB table spaces might not be synchronized.

For a diagram of CHECK LOB syntax and a description of available options, see "Syntax and options of the control statement". For detailed guidance on running this utility, see "Instructions for running CHECK LOB" on page 77.

Output: After successful execution, CHECK LOB resets the CHECK pending (CHKP) and auxiliary warning (AUXW) statuses.

Authorization required: To run this utility, the privilege set of this process must include one of the following authorities:
- STATS privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority

An ID with installation SYSOPR authority can also execute CHECK LOB.

Execution phases of CHECK LOB:

Phase	Description
UTILINIT	Initialization
CHECKLOB	Scans all active pages of the LOB table space
SORT	Sorts four types of records from the CHECKLOB phase; reports four times the number of rows sorted.
REPRTLOB	Examines records that are produced by the CHECKLOB phase and sorted by the SORT phase, and issues error messages
UTILTERM	Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

© Copyright IBM Corp. 1983, 2001

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

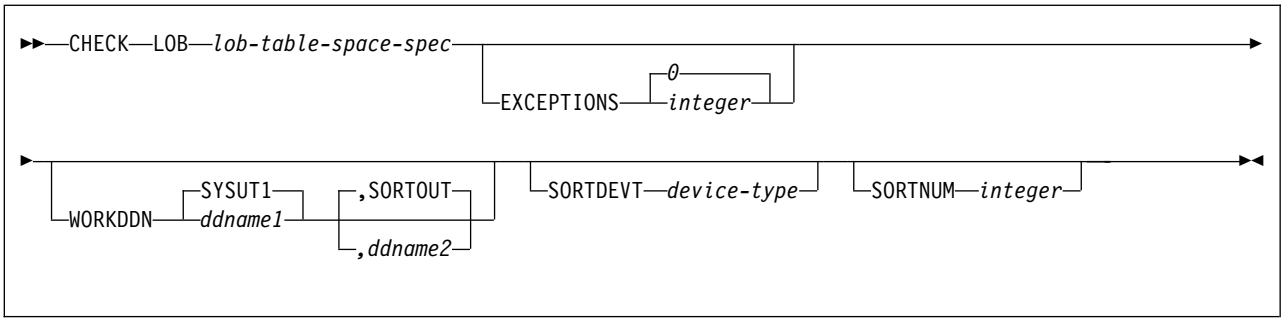

lob-table-space-spec:

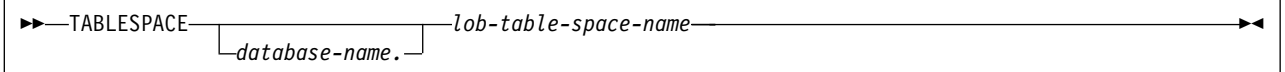

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

LOB Indicates that you are checking a LOB table space for defects.

TABLESPACE *database-name.lob-table-space-name*
Specifies the table space to which the data belongs.

database-name is the name of the database and is optional. The **default** is **DSNDB04**.

lob-table-space-name is the name of the LOB table space.

EXCEPTIONS *integer*
Specifies the maximum number of exceptions, which are reported by messages only. CHECK LOB terminates in the CHECKLOB phase when it reaches the specified number of exceptions.

All defects that are reported by messages are applied to the exception count.

integer is the maximum number of exceptions. The **default** is **0**, which indicates no limit on the number of exceptions.

WORKDDN(*ddname1*,*ddname2***)**
Specifies the DD statements for the temporary work file for sort input and the temporary work file for sort output. A temporary work file for sort input and output is *required.*

ddname1 is the DD name of the temporary work file for sort input.

The **default** is **SYSUT1**.

ddname2 is the DD name of the temporary work file for sort output.

The **default** is **SORTOUT**.

The **WORKDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

SORTDEVT *device-type*

Specifies the device type for temporary data sets that are to be dynamically allocated by DFSORT.

device-type is the device type. *device-type* can be any device type that is acceptable to the DYNALLOC parameter of the SORT or OPTION control statement for DFSORT, as described in *DFSORT Application Programming: Guide.*

If you omit SORTDEVT and a sort is required, you must provide the DD statements that the sort program requires for the temporary data sets.

A TEMPLATE specification does not dynamically allocate sort work data sets. The SORTDEVT keyword controls dynamic allocation of these data sets.

SORTNUM *integer*

Indicates the number of temporary data sets that are to be dynamically allocated by the sort program.

integer is the number of temporary data sets.

If you omit SORTDEVT, SORTNUM is ignored.

If you use SORTDEVT and omit SORTNUM, no value is passed to DFSORT, which then uses its own default.

Instructions for running CHECK LOB

To run CHECK LOB:

1. Read "Before running CHECK LOB" on page 78.
2. Prepare the necessary data sets, as described in "Data sets used by CHECK LOB" on page 78.
3. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for CHECK LOB, see "Sample control statements" on page 80.)
4. Prepare a utility control statement that specifies the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 78. (For a complete description of the syntax and options for CHECK LOB, see "Syntax diagram" on page 76.)
5. Check the compatibility table in "Concurrency and compatibility" on page 79 if you want to run other jobs concurrently on the same target objects.
6. Plan for restarting CHECK LOB if the job doesn't complete, as described in "Terminating or restarting CHECK LOB" on page 79.
7. Run CHECK LOB.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for a description of ways to execute DB2 utilities.

Before running CHECK LOB

You must first recover a LOB table space that is in RECOVER-pending status before running CHECK LOB.

Data sets used by CHECK LOB

Table 10 describes the data sets that CHECK LOB uses. Include statements in your JCL for each required data set and any optional data sets you want to use.

Table 10. Data sets used by CHECK LOB

Data Set	Description	Required?
SYSIN	An input data set containing the utility control statement.	Yes
SYSPRINT	An output data set for messages.	Yes
Work data set	One temporary data set for sort input and sort output. The symbolic name of the DD statement is specified with the WORKDDN option of the utility control statement. The default *ddname* for this data set is SYSUT1. To find the approximate size in bytes of the work data sets, see page 78.	Yes
UTPRINT	Contains messages from DFSORT (usually, SYSOUT or DUMMY).	No

The following object is named in the utility control statement and does not require DD statements in the JCL:

Table space
Object to be checked. This object is named in the CHECK LOB control statement and is accessed through the DB2 catalog.

Defining work data sets: Two sequential data sets, described by the DD statements named in the WORKDDN option, are required during execution of CHECK LOB.

To find the approximate size, in bytes, of the WORKDDN data set:
1. Find the *high allocated page number,* either from the NACTIVEF column of the SYSIBM.SYSTABLESPACE catalog table after running the RUNSTATS utility on the LOB table space, or from information in the VSAM catalog data set.
2. Use the formula *(43 × NACTIVEF × 4).* The resulting value is the approximate size, in bytes, of the work data set required.

Creating the control statement

See “Syntax diagram” on page 76 for CHECK LOB syntax and option descriptions. See “Sample control statements” on page 80 for examples of CHECK LOB usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

“Finding and resolving violations” on page 79
“Resetting CHECK-pending status for a LOB table space” on page 79
“Resolving media failure” on page 79

Finding and resolving violations

CHECK LOB issues message DSNU743I whenever a LOB value is invalid. The violation is identified by:

- The row ID and version number of the LOB
- A reason code for the error
- The page number where the error was found

You can resolve LOB violations by using the UPDATE or DELETE SQL statements to update the LOB column or delete the row associated with the LOB (use the rowid given in message DSNU743I). For more information, see UPDATE or DELETE in Chapter 5 of *DB2 SQL Reference*.

If CHECK LOB issues either message DSNU785I or DSNU787I, it has detected a logical inconsistency within the LOB table space. Contact IBM Support Center for assistance with diagnosing and resolving the problem.

Resetting CHECK-pending status for a LOB table space

If you run CHECK LOB and LOB table space errors are found, the table space is placed in CHECK-pending status.

Complete the following tasks to remove the CHECK-pending status:

1. Correct any defects found in the LOB table space using the REPAIR utility.
2. To reset CHECK-pending or AUXW status, run CHECK LOB again, or run the REPAIR utility.

Resolving media failure

Run CHECK LOB on a LOB table space when a media failure has occurred that leaves LOB pages in the LPL.

Terminating or restarting CHECK LOB

Terminating CHECK LOB

If you terminate CHECK LOB during the CHECKLOB phase, it sets the table space to CHECK-pending status. Otherwise, CHECK LOB resets the CHECK-pending status at the end of the phase if no errors are detected.

For instructions on terminating an online utility, see “Terminating an online utility with the TERM UTILITY command” on page 45.

Restarting CHECK LOB

You can restart a CHECK LOB utility job, but it starts from the beginning again.

For instructions on restarting a utility job, see “Restarting an online utility” on page 46.

Concurrency and compatibility

Claims and drains: Table 11 shows the claim classes in which CHECK LOB claims and drains and any restrictive state the utility sets on the target object.

Table 11. Claim classes on a LOB table space and index on the auxiliary table for CHECK LOB operations. Use of claims and drains; restrictive states set on the target object.

TARGET OBJECTS	CHECK LOB
LOB table space	DW/UTRO
Index on the auxiliary table	DW/UTRO

Table 11. Claim classes on a LOB table space and index on the auxiliary table for CHECK LOB operations (continued). Use of claims and drains; restrictive states set on the target object.

TARGET OBJECTS	CHECK LOB

Legend:
- DW: Drain the write claim class, concurrent access for SQL readers
- UTRO: Utility restrictive state, read only access allowed

Compatibility: Any SQL operation or other online utility that attempts to update the same LOB table space is incompatible.

Sample control statements

Example: Checking a LOB table space. Check the table space TLIQUG02 in database DBIQUG01 for structural defects or invalid LOB values.

```
//STEP1    EXEC DSNUPROC,UID='IUIQU2UG.CHECKL',
//         UTPROC='',
//         SYSTEM='V71A'
//SYSERR   DD DSN=IUIQU2UG.STEP1.SYSERR,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD DSN=IUIQU2UG.CHECKL.SYSUT1,UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND),
//         DISP=(MOD,DELETE,CATLG)
//SYSREC   DD DSN=CUST.FM.CSFT320.DATA,DISP=SHR,UNIT=SYSDA,
//         VOL=SER=123456
//SORTOUT  DD DSN=IUIQU2UG.STEP1.SORTOUT,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
CHECK LOB TABLESPACE DBIQUG01.TLIQUG02
      EXCEPTIONS 3 WORKDDN SYSUT1,SORTOUT SORTDEVT SYSDA
      SORTNUM 4
```

Chapter 9. COPY

The COPY online utility creates up to four image copies of any of the following objects:

- Table space
- Table space partition
- Data set of a linear table space
- Index space
- Index space partition

There are two types of image copies:

1. A *full image copy* is a copy of all pages in a table space, partition, data set, or index space.
2. An *incremental image copy* is a copy only of pages that have been modified since the last use of the COPY utility.

The copies are used by the RECOVER utility when recovering a table space or
| index space to the most recent time or to a previous time. Copies can also be used
| by MERGECOPY, RECOVER, COPYTOCOPY and UNLOAD.

You can copy a list of objects in parallel to improve performance. Specifying a list of objects along with the SHRLEVEL REFERENCE option creates a single recovery point for that list of objects. Specifying the PARALLEL keyword allows you to copy a list of objects in parallel, rather than serially. This provides a performance advantage, and allows the RECOVER utility to process the logs for all table spaces and index spaces in a single pass.

The COPY utility uses 3 threads for single object COPY job. To calculate the
number of threads needed when you specify the PARALLEL keyword, use the
formula (n*2+1), where n is the number of objects that should be processed in
parallel, regardless of the total number of objects in the list. If you do not use the
PARALLEL keyword, then n is 1 and COPY uses 3 threads for a single object copy
job.

For a diagram of COPY syntax and a description of available options, see “Syntax and options of the control statement” on page 82. For detailed guidance on running this utility, see “Instructions for running COPY” on page 89.

Output: Output from the COPY utility consists of:

- Up to four sequential data sets containing the image copy.
- Rows in the SYSIBM.SYSCOPY catalog table that describe the image copy data sets available to the RECOVER utility. It is your installation’s responsibility to ensure that these data sets are available if the RECOVER utility requests them.
- If you specify the CHANGELIMIT option, a report on the change status of the table space.

The COPY-pending status is off for table spaces if the copy was a full image copy. However, DB2 does not reset the COPY-pending status if you COPY a single piece of a multi-piece linear data set. If you COPY a single table space partition, DB2 resets the COPY-pending status only for the copied partition and not the whole table space. DB2 resets the informational COPY-pending status (ICOPY) after you copy an index space or index.

© Copyright IBM Corp. 1983, 2001

Related information: See Part 4 (Volume 1) of *DB2 Administration Guide* for uses of COPY in the context of planning for database recovery. For information about creating copies inline during LOAD, see “Using inline COPY with LOAD” on page 194. You can also create inline copies during REORG; see “Using inline COPY with REORG TABLESPACE” on page 373 for more information.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- IMAGCOPY privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority

An ID with installation SYSOPR authority can also execute COPY, but only on a table space in the DSNDB01 or DSNDB06 database.

The batch user ID that invokes COPY with the CONCURRENT option must provide the necessary authority to execute the DFDSS DUMP command.

Execution phases of COPY: The COPY utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup
REPORT	Reporting for CHANGELIMIT option
COPY	Copying
UTILTERM	Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

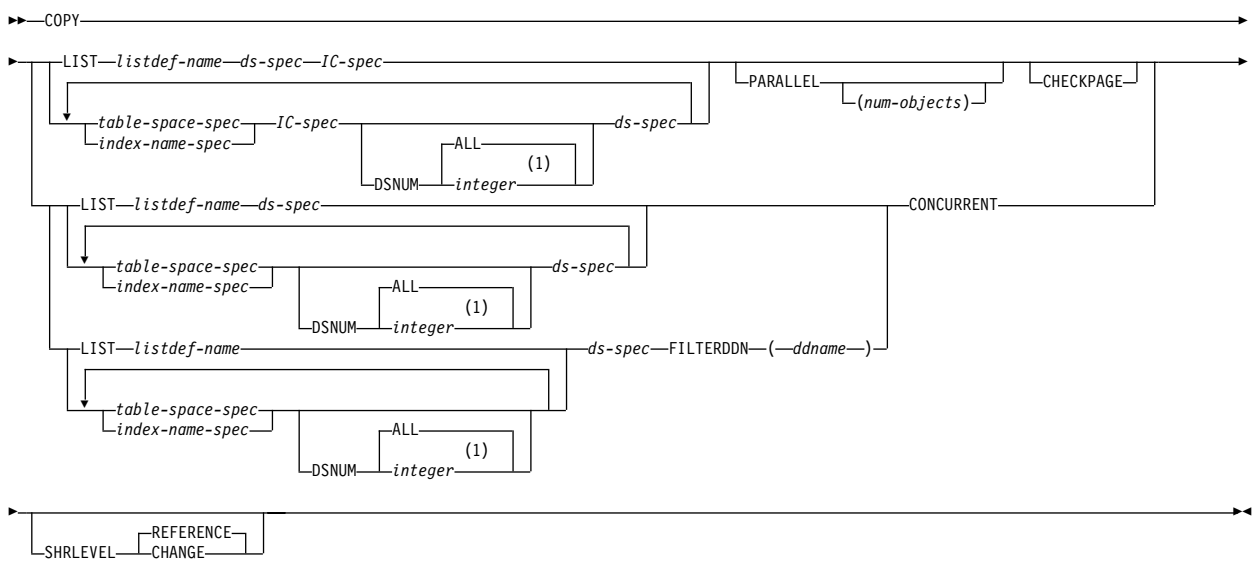

Notes:

1 Not valid for nonpartitioning indexes.

IC-spec:

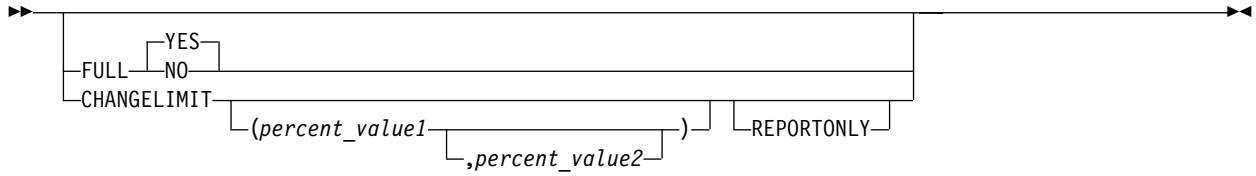

table-space-spec:

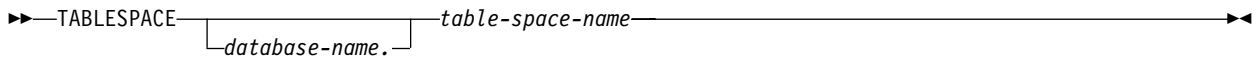

index-name-spec:

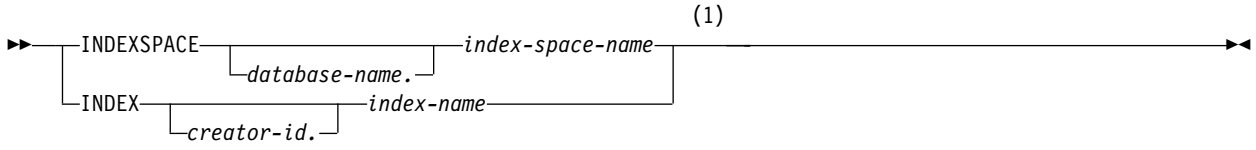

Notes:

1 INDEXSPACE is the preferred specification.

ds-spec:

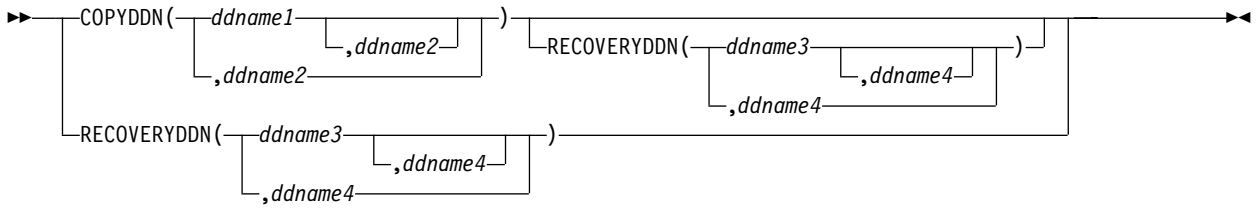

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

LIST *listdef-name*

Specifies the name of a previously defined LISTDEF list name. LIST specifies one LIST keyword for each COPY control statement. Do not specify LIST with either the INDEX keywords or the TABLESPACE, utility list manager will invoke COPY once for the entire list. For more information about LISTDEF specifications, see "Generating object lists using LISTDEF" on page 144.

TABLESPACE

Specifies the table space (and, optionally, the database it belongs to) that is to be copied.

database-name

Is the name of the database the table space belongs to.

The **default** is **DSNDB04**.

table-space-name

Is the name of the table space to be copied.

Specify the DSNDB01.SYSUTILX, DSNDB06.SYSCOPY, or DSNDB01.SYSLGRNX table space by itself in a single COPY statement, or only with indexes over the table space that were defined with the COPY YES attribute.

INDEXSPACE *database-name.index-space-name*
Specifies the qualified name of the index space that is to be copied; the name is obtained from the SYSIBM.SYSINDEXES table. The index space specified must be defined with the COPY YES attribute.

database-name
Optionally specifies the name of the database the index space belongs to.

The **default** is **DSNDB04**.

index-space-name
Specifies the name of the index space to be copied.

INDEX *creator-id.index-name*
Specifies the index to be copied.

creator-id
Optionally specifies the creator of the index.

The **default** is the user identifier for the utility.

index-name
Specifies the name of the index that is to be copied.

DSNUM For a table space, identifies a partition or data set within the table space to be copied; or it copies the entire table space. For an index space, identifies a partition to be copied; or it copies the entire index space.

If a data set of a nonpartitioned table space is in the COPY-pending status, you must copy the entire table space.

ALL Copies the entire table space or index space.

The **default** is **ALL**.

You must use ALL for a non-partitioning index.

integer
Is the number of a partition or data set to be copied.

An integer value is not valid for non-partitioning indexes.

For a partitioned table space or index space, the integer is its partition number. The maximum is 254.

For a nonpartitioned table space, find the integer at the end of the data set name as cataloged in the VSAM catalog. The data set name has this format:

catname.DSNDB*x*.*dbname*.*spacename*.*y*0001.A*nnn*

where:

catname	The VSAM catalog name or alias
x	C or D
dbname	The database name
spacename	The table space or index space name
y	I or J
nnn	The data set integer.

If image copies are taken by data set (rather than by table space), then RECOVER, MERGECOPY or COPYTOCOPY must use the copies by data set. For a nonpartitioned table space, if image

copies are taken by data set and you run MODIFY RECOVERY with DSNUM ALL, then the table space is placed in COPY-pending status if a full image copy of the entire table space does not exist.

COPYDDN *ddname1,ddname2*

Specifies a DD name or a TEMPLATE name for the primary (*ddname1*) and backup (*ddname2*) copied data sets for the image copy at the local site.

ddname is the DD name.

The **default** is **SYSCOPY** for the primary copy. You can only use the default for one object in the list.

If you use the CHANGELIMIT REPORTONLY option, you may use a DD DUMMY card when you specify the SYSCOPY output data set. This card prevents a data set from being allocated and opened.

It is recommended that you catalog all of your image copy data sets.

You cannot have duplicate image copy data sets. If the DD statement identifies a noncataloged data set with the same name, volume serial, and file sequence number as one already recorded in SYSIBM.SYSCOPY, a message is issued and no copy is made. If it identifies a cataloged data set with only the same name, no copy is made. For cataloged image copy data sets, CATLG must be specified for the normal termination disposition in the DD statement; for example, DISP=(MOD,CATLG,CATL). The DSVOLSER field of the SYSIBM.SYSCOPY entry will be blank.

When the image copy data set is going to a tape volume, the VOL=SER parameter should be specified on the DD statement.

If you use the CONCURRENT and FILTERDDN options, make sure the size of the copy data set is large enough to include all of the objects in the list.

The **COPYDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

RECOVERYDDN *ddname3,ddname4*

Specifies a DD name or a TEMPLATE name for the primary (*ddname3*) and backup (*ddname4*) copied data sets for the image copy at the recovery site.

ddname is the DD name.

You cannot have duplicate image copy data sets. The same rules apply for RECOVERYDDN as for COPYDDN.

If you use the CONCURRENT and FILTERDDN options, make sure the size of the copy data set is large enough to include all of the objects in the list.

The **RECOVERYDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is

both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

PARALLEL Specifies the maximum number of objects in the list that should be processed in parallel. If you omit this keyword, the list is not processed in parallel.

(*num-objects***)**
Specifies the number of objects in the list that should be processed in parallel. This value can be adjusted to a smaller value if COPY encounters storage constraints.

If you specify 0 or do not specify a value for *num-objects*, COPY determines the optimal number of objects to process in parallel.

See “Copying a list of objects” on page 95.

CHECKPAGE If specified, checks each page in the table space or index space for validity. The validity checking operates on one page at a time and does not include any cross-page checking. If an error is found, a message is issued describing the type of error. If more than one error exists in a given page, only the first error is identified. COPY will continue checking the remaining pages in the table space or index space after an error is found.

FULL Makes either a full or an incremental image copy.

YES Makes a full image copy.

Making a full image copy resets the COPY-pending status for the table space or index, or for the partition if you specified DSNUM.

The **default** is **YES**.

NO Makes only an incremental image copy. Only changes since the last image copy are copied.

NO is not valid for indexes.

Incremental image copies are not allowed in the following situations:

- The last full image copy of the table space was taken with the CONCURRENT option.
- No full image copies exist for the table space or data set being copied.
- After a successful LOAD or REORG operation, unless an inline copy was made during the LOAD or REORG.
- The table space you specify is one of the following: DSNDB01.DBD01, DSNDB01.SYSUTILX, or DSNDB06.SYSCOPY.

COPY automatically takes a full image copy of a table space if you specify FULL NO when an incremental image copy is not allowed.

FILTERDDN *ddname*
Specifies the DD statement for the filter data set to be used, if desired, by COPY with the CONCURRENT option. COPY uses this

data set to automatically build a list of table spaces to be copied by DFSMSdss with one DFSMSdss ″DUMP″ statement.

If FILTERDDN is specified, the SYSIBM.SYSCOPY records for all objects in the list will have the same data set name.

ddname is the DD name.

The **FILTERDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

SHRLEVEL Indicates whether other programs can access or update the table space or index while COPY is running.

REFERENCE
Allows read-only access by other programs.

The **default** is **REFERENCE**.

CHANGE
Allows other programs to change the table space or index space.

When you specify SHRLEVEL CHANGE, uncommitted data might be copied. Image copies taken using SHRLEVEL CHANGE are not recommended for use with RECOVER TOCOPY.

SHRLEVEL CHANGE is not allowed when you use DFSMS Concurrent Copy for table spaces having a page size greater than 4KB.

CONCURRENT
Executes DFSMS concurrent copy to make the full image copy. The image copy is recorded in SYSCOPY with ICTYPE=F and STYPE=C.

If the SYSPRINT DD card points to a data set, you must use a DSSPRINT DD card.

When SHRLEVEL(REFERENCE) is specified, an ICTYPE=Q record is placed into the SYSCOPY table after the object has been quiesced. If COPY fails, then this record remains in SYSCOPY. When COPY is successful, then this ICTYPE=Q record is replaced with the ICTYPE=F record.

For table spaces with a 32 KB page size, you must run the job with the SHRLEVEL REFERENCE (default) option when using the CONCURRENT option. Otherwise, the job is terminated, and message DSNU423I is issued.

CHANGELIMIT
Specifies the percent limit of changed pages in the table space, partition, or data set when an incremental or full image copy should be taken.

percent_value1
Specifies a value in the CHANGELIMIT range. *percent_value1*

must be an integer or decimal value from 0.0 to 100.0. You do not need to specify leading zeroes, and the decimal point is not required when specifying a whole integer. Specify a maximum ofone decimal place for a decimal value (for example, .5).

percent_value2
Specifies the second value in the CHANGELIMIT range. *percent_value2* must be an integer or decimal value from 0.0 to 100.0. You do not need to specify leading zeroes, and the decimal point is not required when specifying a whole integer. Specify a maximum of one decimal place for a decimal value (for example, .5).

COPY CHANGELIMIT accepts values in any order.

If only one value is specified, COPY CHANGELIMIT:

- Creates an incremental image copy if the percentage of changed pages is greater than 0 and less than *percent_value1*.
- Creates a full image copy if the percentage of change pages is greater than or equal to *percent_value1*, or if CHANGELIMIT(0) is specified.
- Does not create an image copy if no pages have changed, unless CHANGELIMIT(0) is specified.

If two values are specified, COPY CHANGELIMIT:

- Creates an incremental image copy if the percentage of changed pages is greater than the lowest value specified and less than the highest value specified.
- Creates a full image copy if the percentage of changed pages is equal to or greater than the highest value specified.
- Does not create an image copy if the percentage of changed pages is less than or equal to the lowest value specified.
- If both values are equal, creates a full image copy if the percentage of changed pages is equal to or greater than the value specified.

The **default** values are **(1,10)**.

You cannot use the CHANGELIMIT option for a table space or partition defined with TRACKMOD NO. If you change the TRACKMOD option from NO to YES, you must take an image copy before you can use the CHANGELIMIT option. For nonpartitioned table spaces, you must copy the entire table space to allow future CHANGELIMIT requests.

REPORTONLY
Specifies that image copy information is displayed. If you specify the REPORTONLY option, then only image copy information is displayed. Image copies are not taken, only recommended.

Instructions for running COPY

To run COPY, you must:

1. Read “Before running COPY” on page 90 in this chapter.

2. Prepare the necessary data sets, as described in “Data sets used by COPY”.
3. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for COPY, see “Sample control statements” on page 104.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 92. (For a complete description of the syntax and options for COPY, see “Syntax and options of the control statement” on page 82.)
5. Check the compatibility table in “Concurrency and compatibility” on page 102 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the COPY job doesn’t complete, as described in “Terminating or restarting COPY” on page 101.
7. Run COPY.

See “Chapter 3. Invoking DB2 online utilities” on page 23 for an explanation of ways to execute DB2 utilities.

Before running COPY

Checking table space status: You cannot copy a table space that is in the CHECK-pending or RECOVER-pending status. See “Resetting RECOVER-pending or REBUILD-pending status” on page 297 for information about resetting these statuses.

Resetting COPY-pending status: If a table space is in COPY-pending status, or an index is in informational COPY-pending status, you can reset the status only by taking a full image copy of the entire table space, all partitions of the table space, or the index space. When you make an image copy of a partition, the COPY-pending status of the partition is reset. If a nonpartitioned table space is in COPY-pending status, you can reset the status only by taking a full image copy of the entire table space, and not of each data set.

Data sets used by COPY

Table 12 describes the data sets required for COPY. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 12. Data sets used by COPY

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
DSSPRINT	Output data set for messages; required when CONCURRENT copy is used and the SYSPRINT DD card points to a data set.	No
Filter	A single data set DB2 uses when you specify the FILTERDDN option in the utility control statement; contains a list of VSAM data set names built by DB2, and is used during COPY when the CONCURRENT and FILTERDDN options are specified.	Yes[1]

Table 12. Data sets used by COPY (continued)

Data Set	Description	Required?
Copies	From one to four output data sets to contain the resulting image copy data sets. Their DD names are specified with the COPYDDN and RECOVERYDDN options of the utility control statement. The default is one copy, in the data set described by the SYSCOPY DD statement.	Yes

Note: [1] Required if you specify the FILTERDDN option.

The following objects are named in the utility control statement and do not require DD cards in the JCL:

Table space or Index space
Object to be copied. It is named in the COPY control statement and is accessed through the DB2 catalog. (If you want to copy only certain data sets in a table space, you must use the DSNUM option in the control statement.)

DB2 catalog objects
COPY records each copy in the DB2 catalog table SYSIBM.SYSCOPY.

Output data set size: Image copies are written to sequential non-VSAM data sets. To find the approximate size of the image copy data set for a table space, in bytes, you can execute COPY with the CHANGELIMIT REPORTONLY option, or use the following procedure:

1. Find the *high allocated page number,* either from the NACTIVE column of SYSIBM.SYSTABLESPACE after running the RUNSTATS utility, or from information in the VSAM catalog data set.
2. Multiply the high allocated page number by the page size.

Filter data set size: Use the formula *(240 + (80 × n))* to determine the approximate FILTER data set size required, in bytes, where *n* = the number of objects specified in the COPY control statement.

JCL parameters: You can specify a block size for the output by using the BLKSIZE parameter on the DD statement for the output data set. Valid block sizes are multiples of 4096 bytes. You can increase the buffering with the BUFNO parameter; for example, you might specify BUFNO=30. See also “Data sets used by online utilities” on page 24 for information about using BUFNO.

Cataloging image copies: To catalog your image copy data sets, use the DISP=(MOD,CATLG,CATLG) parameter in the DD statement named by the COPYDDN option. After the image copy is taken, the DSVOLSER column of the row inserted into SYSIBM.SYSCOPY contains blanks.

Duplicate image copy data sets are not allowed. If there is a cataloged data set already recorded in SYSCOPY with the same name as the new image copy data set, a message is issued and the copy is not made.

When RECOVER locates the entry in SYSCOPY, it uses the MVS catalog to allocate the required data set. *If you have uncataloged the data set, the allocation fails.* In that case, the recovery can still go forward; RECOVER searches for a previous image copy. But even if it finds one, it must use correspondingly more of

the log to recover. It is to your benefit, and it is your responsibility, to keep the MVS catalog consistent with SYSIBM.SYSCOPY about existing image copy data sets.

Creating the control statement

See “Syntax diagram” on page 82 for COPY syntax and option descriptions. See “Sample control statements” on page 104 for examples of COPY usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

“Making full image copies”
“Making incremental image copies” on page 93
“Making multiple image copies” on page 93
“Copying partitions or data sets in separate jobs” on page 95
“Copying a list of objects” on page 95
“Using more than one COPY statement” on page 97
“Copying segmented table spaces” on page 97
“Using DFSMS concurrent copy” on page 97
“Specifying conditional image copies” on page 98
“Preparing for recovery” on page 99
“Improving performance” on page 100

Making full image copies

You can make a full image copy of any of the following objects:

- Table space
- Table space partition
- Data set of a linear table space
- Index space
- Index space partition

The following statement makes a full image copy of the DSN8S71E table space in database DSN8D71A:

```
COPY TABLESPACE DSN8D71A.DSN8S71E
```

The COPY utility writes pages from the table space or index space to the output data sets. The JCL for the utility job must include DD statements for the data sets. If the object consists of multiple data sets and all are copied in one run, the copies reside in one physical sequential output data set.

Image copies should be made either by entire page set or by partition, but not by both. We recommend taking a full image copy after CREATE or LOAD operations for a new object that is populated, after a REORG operation for an existing object, and after LOAD RESUME of an existing object.

We recommend copying the indexes over a table space whenever a full copy of the table space is taken. More frequent index copies decrease the number of log records to be applied during recovery. At minimum, you should copy an index when it is placed in informational COPY-pending (ICOPY) status. For more information about the ICOPY status, see “Appendix C. Resetting an advisory or restrictive status” on page 677.

If you create an inline copy during LOAD or REORG, you do not need to execute a separate COPY job for the table space. If you do not create an inline copy, and if the LOG option was NO, the COPY-pending status is set for the table space. A full image copy must be made for any subsequent recovery of the data. An incremental image copy is not allowed.

If the LOG option was YES, the COPY-pending status is not set. However, your next image copy must be a full image copy. Again, an incremental image copy is not allowed.

The COPY utility will automatically take a full image copy of a table space if you attempt to take an incremental image copy when it is not allowed.

The catalog table SYSIBM.SYSCOPY and the directory tables SYSIBM.SYSUTILX and SYSIBM.SYSLGRNX record information from the COPY utility. Copying the catalog table or the directories can lock out separate COPY jobs that are running simultaneously; therefore, it is most efficient to defer copying the catalog table or directories until the other copy jobs have completed. However, if you must copy other objects while another COPY job processes catalog tables or directories, specify SHRLEVEL (CHANGE) for the copies of the catalog and directory tables.

Making incremental image copies

An incremental image copy is a copy of the pages that have been changed since the last full or incremental image copy. You cannot take an incremental image copy of an index space. You can make an incremental image copy of a table space if:

- A full image copy of the table space exists
- The COPY-pending status is not on for that table space
- The last copy was taken without the CONCURRENT option

Copy by partition or data set: You can make an incremental image copy by partition or data set (specified by DSNUM) if a full image copy of the table space exists, or if a full image copy of the same partition or data set exists and the COPY-pending status is not on for the table space or partition. Moreover, the full image copy must have been made after the most recent application to the table space of CREATE, REORG or LOAD, or it must be an inline copy made during the most recent application of LOAD or REORG.

Sample control statement: To specify an incremental image copy, use FULL NO on the COPY statement, as in this example:

```
COPY TABLESPACE DSN8D71A.DSN8S71E
  FULL NO
  SHRLEVEL CHANGE
```

Performance advantage: An incremental image copy generally does not require a complete scan of the table space, with two exceptions:

- The table space was defined with the TRACKMOD NO option.
- You are taking the first copy after you altered a table space to TRACKMOD YES.

Space maps in each table space indicate, for each page, whether it has changed since the last image copy. Therefore, making an incremental copy can be significantly faster than making a full copy if the table space was defined with the TRACKMOD YES option. Incremental image copies of a table space that was defined with TRACKMOD NO will still save space, at some performance cost.

Restrictions: You cannot make incremental copies of the DSNDB01.DBD01 and copies of table space DSNDB01.SYSUTILX in the directory, or DSNDB06.SYSCOPY in the catalog. For those objects, COPY always makes a full image copy and places the SYSCOPY record in the log.

Making multiple image copies

You can use a single invocation of the COPY utility to create up to four exact copies of any of the following objects:

- Table space

- Table space partition
- Data set of a linear table space
- Index space
- Index space partition

Two copies can be made for use on the local DB2 system (installed with the option LOCALSITE), and two more for offsite recovery (on any system installed with the option RECOVERYSITE). All copies are identical, and are produced at the same time from one invocation of COPY.

The ICBACKUP column in SYSIBM.SYSCOPY specifies whether the image copy data set is for the local or recovery system, and whether the image copy data set is for the primary copied data set or the backup copied data set. The ICUNIT column in SYSIBM.SYSCOPY specifies whether the image copy data set is on tape or DASD.

Remote site recovery: In preparation for remote site recovery, system and application libraries and the DB2 catalog and directory are assumed to be identical at the local site and recovery site. You can regularly transport copies of archive logs and database data sets to a safe location to keep data for remote site recovery current. This information can be kept on tape until needed.

Naming the data sets for the copies: The option COPYDDN of COPY names the output data sets that receive copies for local use. The option RECOVERYDDN of COPY names the output data sets that receive copies intended for remote site recovery. The options have these formats:

```
COPYDDN (ddname1,ddname2)

RECOVERYDDN (ddname3,ddname4)
```

The ddnames for the primary output data sets are *ddname1* and *ddname3*. The ddnames for the backup output data sets are *ddname2* and *ddname4*.

Sample control statement: The following statement makes four full image copies of the table space DSN8S71E in database DSN8D71A, using LOCALDD1 and LOCALDD2 as ddnames for the primary and backup copies used on the local system and RECOVDD1 and RECOVDD2 as ddnames for the primary and backup copies for remote site recovery:

```
COPY TABLESPACE DSN8D71A.DSN8S71E
  COPYDDN (LOCALDD1,LOCALDD2)
  RECOVERYDDN (RECOVDD1,RECOVDD2)
```

You do not have to make copies for local use and for remote site recovery at the same time. COPY allows you to use either option COPYDDN or option RECOVERYDDN without the other. If you make copies for local use more often than copies for remote site recovery, then a remote site recovery might work from an older copy, and more of the log, than a local recovery; hence, it would take longer. But, in your plans for remote site recovery, that difference might be acceptable. You can also use MERGECOPY RECOVERYDDN to create recovery site full copies, and merge local incremental copies into new recovery site full copies. Or, you can use COPYTOCOPY to create the missing image copies. See "Chapter 10. COPYTOCOPY" on page 109 for more information.

Making multiple incremental image copies: DB2 cannot make incremental image copies if:

- The incremental image copy is requested only for a site other than the current site (the local site from which the request is made).

- Incremental image copies are requested for both sites, but the most recent full image copy was made for only one site.
- Incremental image copies are requested for both sites and the most recent full image copies were made for both sites, but between the most recent full image copy and current request, incremental image copies were made for the current site only.

If you attempt to make incremental image copies under any of these conditions, COPY terminates with return code 8, does not take the image copy or update the SYSCOPY table, and issues this message:

```
DSNU404I  csect-name LOCAL SITE AND RECOVERY SITE INCREMENTAL
          IMAGE COPIES ARE NOT SYNCHRONIZED
```

To proceed, and still keep the two sets of data synchronized, take another full image copy of the table space for both sites, or change your request to make an incremental image copy only for the site at which you are working.

DB2 cannot make an incremental image copy if the object being copied is an index or index space.

Maintaining copy consistency: Make full image copies for both the local and recovery sites:

- If a table space is in COPY-pending status.
- After a LOAD or REORG procedure that did not create an inline copy.
- If an index is in the informational COPY-pending status.

This action helps to insure correct recovery for both local and recovery sites. If the requested full image copy is for one site only, but the history shows that copies were made previously for both sites, COPY continues to process the image copy and issues the following warning message:

```
DSNU406I  FULL IMAGE COPY SHOULD BE TAKEN FOR BOTH LOCAL SITE AND
          RECOVERY SITE.
```

The COPY-pending status of a table space is not changed for the other site when you make multiple image copies at the current site for that other site. For example, if a table space is in COPY-pending status at the current site, and you make copies from there for the other site only, the COPY-pending status will still be on when you bring up the system at that other site.

Copying partitions or data sets in separate jobs

If you have a partitioned table space or partitioning index, you can copy the partitions independently in separate simultaneous jobs. This can reduce the time it takes to create an image copy of the total table space.

If a nonpartitioned table space consists of more than one data set, you can copy several or all of the data sets independently in separate jobs. To do so, run simultaneous COPY jobs (one job for each data set) and specify SHRLEVEL CHANGE on each.

However, creating copies simultaneously will not provide you with a consistent recovery point unless you follow up with a QUIESCE of the table space.

Copying a list of objects

Within a single COPY statement, the COPY utility allows you to process a list containing any of the following objects:

- Table space
- Table space partition

- Data set of a linear table space
- Index space
- Index space partition

Specifying objects in a list is useful for copying a complete set of referentially related table spaces before running QUIESCE. Consider the following information when taking an image copy for a list of objects:

- DB2 copies table spaces and index spaces in the list one at a time, in the specified order, unless you invoke parallelism by specifying the PARALLEL keyword.
- Each table space in the list with a CHANGELIMIT specification will have a REPORT phase, so the phase will switch between REPORT and COPY while processing the list.
- If processing completes successfully, any COPY-pending status on the table spaces and informational COPY-pending status on the indexes will be reset.
- Using the SHRLEVEL(REFERENCE) option:

  Drains the write claim class on each table space and index in the UTILINIT phase, which is held for the duration of utility processing.

  Utility processing inserts SYSIBM.SYSCOPY rows for all of the objects in the list at the same time, after all of the objects have been copied.

  All objects in the list will have identical RBA or LRSN values for the START_RBA column for the SYSIBM.SYSCOPY rows: the current LRSN at the end of the COPY phase.
- Using the SHRLEVEL(CHANGE) option:

  Claims the read class for each table space and index space. The claim initializes before the copy of the object starts, and releases when the copy on the object completes.

  Utility processing inserts a SYSIBM.SYSCOPY row for objects in the list when the copy of each object is complete.

  Objects in the list will have different LRSN values for the START_RBA column for the SYSIBM.SYSCOPY rows: the current RBA or LRSN at the start of copy processing for that object.

When you specify the PARALLEL keyword, DB2 supports parallelism for image copies on DASD devices. To calculate the number of threads needed when you specify the PARALLEL keyword, use the formula (n*2+1), where n is the number of objects that should be processed in parallel, regardless of the total number of objects in the list. If you do not use the PARALLEL keyword, then n is 1 and COPY uses 3 threads for a single object copy job.

If COPY encounters a tape volume in the list, processing of remaining objects pauses until the tape object has completed, then parallel processing resumes.

The following table spaces cannot be included in a list of table spaces and must be specified as a single object:

DSNDB01.SYSUTILX
DSNDB06.SYSCOPY
DSNDB01.SYSLGRNX

The only exceptions to this restriction are the indexes over these table spaces that were defined with the COPY YES attribute. You can specify such indexes along with the appropriate table space.

Using more than one COPY statement

You can use more than one control statement for COPY in one DB2 utility job step. After each COPY statement has executed successfully:

- A row referring to the image copy is recorded in SYSIBM.SYSCOPY table.
- The image copy data set is valid and available for RECOVER, MERGECOPY, COPYTOCOPY and UNLOAD.

If a job step containing more than one COPY statement abends, **do not use TERM UTILITY.** Restart the job from the last commit point using RESTART instead. Terminating COPY in this case creates inconsistencies between the ICF catalog and DB2 catalogs.

Copying segmented table spaces

COPY distinguishes between segmented and nonsegmented table spaces. If you specify a segmented table space, COPY locates empty and unformatted data pages in the table space and does not copy them.

Using DFSMS concurrent copy

You might be able to gain improved availability by using the Concurrent Copy function of Data Facility Storage Management Subsystem (DFSMS). You can subsequently run the RECOVER utility to restore those image copies and apply the necessary log records to them to complete recovery.

The CONCURRENT option invokes DFSMS concurrent copy. The COPY utility records the resulting DFSMS concurrent copies in the catalog table SYSIBM.SYSCOPY with ICTYPE=F and STYPE=C.

To obtain a consistent offline backup:

1. Start the DB2 objects being backed up for read-only access by issuing the following command:

   ```
   -START DATABASE(database name) SPACENAM(tablespace-name) ACCESS(RO)
   ```

   This is necessary to ensure that no updates to data occur during this procedure.
2. Run QUIESCE with the WRITE(YES) option to quiesce all DB2 objects being backed up.
3. Back up the DB2 data sets if the QUIESCE utility completes successfully.
4. Issue the following command to allow transactions to access the data:

   ```
   -START DATABASE(database name) SPACENAM(tablespace-name)
   ```

If you use the CONCURRENT option:

- You must supply either a COPYDDN ddname, a RECOVERYDDN ddname, or both.
- You can set the disposition to DISP=(MOD,CATLG,CATLG) if you specify the new data set for the image copy on a scratch volume (a specific volume serial number is not specified). You must set the disposition to DISP=(NEW,CATLG,CATLG) if you specify a specific volume serial number for the new image copy data set.
- If you are restarting COPY, specify DISP=(MOD,CATLG,CATLG) or DISP=(NEW,CATLG,CATLG) for the COPYDDN and RECOVERYDDN data sets. The DFSMSdss DUMP command does not support appending to an existing data set. Therefore, the COPY utility covers any DISP=MOD data sets to DISP=OLD before invoking DFSMS.
- If the SYSPRINT DD card points to a data set, you must use a DSSPRINT DD card.

- You must use the SHRLEVEL REFERENCE option for table spaces with a 8KB, 16KB, or 32KB page size.

Restrictions on using DFSMS concurrent copy: You cannot use a copy made with DFSMS concurrent copy with the PAGE or ERROR RANGE options. If you specify PAGE or ERROR RANGE, RECOVER bypasses any concurrent copy records when searching the SYSCOPY table for a recoverable point.

You cannot use the CONCURRENT option with SHRLEVEL CHANGE on a table space with 8 KB, 16 KB, or 32 KB page size.

Also, you cannot run the following DB2 stand-alone utilities on copies made by DFSMS concurrent copy:

DSN1COMP
DSN1COPY
DSN1PRNT

You cannot execute the CONCURRENT option from the DB2I Utilities panel or from the DSNU TSO CLIST command.

Requirements for using DFSMS concurrent copy: To use COPY to take DFSMS concurrent copies, you must have the following hardware and software:

- OS/390 Release 3
- 3990 model 3 or 3990 model 6 controller at the extended platform attached to the DASD. A COPY job fails if one or more of the table spaces names is on DASD that does not have the controller.

Table space availability: If you specify COPY SHRLEVEL REFERENCE with the CONCURRENT option, and if you want to copy all of the data sets for a list of table spaces to the same output device, specify FILTERDDN in your COPY statement to improve table space availability. In this scenario, specifying COPY without the FILTERDDN option forces DFSMS to process the list of table spaces sequentially, which might limit the availability of some of the table spaces being copied.

Specifying conditional image copies

Use the CHANGELIMIT option of the COPY utility to specify conditional image copies. You can use it to get a report of image copy information about a table space, or you can let DB2 decide whether to take an image copy based on this information.

You cannot use the CHANGELIMIT option for a table space or partition defined with TRACKMOD NO. If you change the TRACKMOD option from NO to YES, you must take an image copy before you can use the CHANGELIMIT option. When you change the TRACKMOD option from NO to YES for a linear table space, you must take a full image copy using DSNUM ALL before you can copy using the CHANGELIMIT option.

Obtaining image copy information about a table space: When you specify COPY CHANGELIMIT REPORTONLY, COPY reports image copy information for the table space and recommends the type of copy, if any, to take. The report includes:

- The total number of pages in the table space. This value is the number of pages copied if a full image copy is taken.
- The number of empty pages, if the table space is segmented.

- The number of changed pages. This value is the number of pages copied if an incremental image copy is taken.
- The percentage of changed pages.
- The type of image copy recommended.

Adding conditional code to your COPY job: You can add conditional code to your jobs so that an incremental or full image copy, or some other step, is performed depending on how much the table space has changed. For example, you can add a conditional MERGECOPY step to create a new full image copy if your COPY job took an incremental copy. COPY CHANGELIMIT uses the following return codes to indicate the degree that a table space or list of table spaces has changed:

1 (informational)	If no CHANGELIMIT was met.
2 (informational)	If the percent of changed pages is greater than the low CHANGELIMIT and less than the high CHANGELIMIT value.
3 (informational)	If the percent of changed pages is greater than or equal to the high CHANGELIMIT value.

If you specify multiple copy statements in one job step, that job step will report the highest return code from all of the imbedded statements. Basically, the statement with the highest percentage of changed pages determines the return code and the recommended action against the entire list of COPY statements contained in the subsequent job step.

Using conditional copy with generation data groups (GDGs): When you use generation data groups (GDGs) and need to make an incremental image copy, take the following steps to prevent creating an empty image copy:

1. Include in your job a first step in which you run COPY with CHANGELIMIT REPORTONLY. Set the SYSCOPY DD card to DD DUMMY so no output data set is allocated.
2. Add a conditional JCL statement to examine the return code from the COPY CHANGELIMIT REPORTONLY step.
3. Add a second COPY step without CHANGELIMIT REPORTONLY to copy the table space or table space list based on the return code from the first step.

Preparing for recovery

If you are taking incremental copies, if you have recently run REORG or LOAD, or if you plan to recover a LOB table space, read the following topics pertaining to recovery.

Using incremental copies: The RECOVER TABLESPACE utility merges all incremental image copies since the last full image copy, and must have all the image copies available at the same time. If there is any likelihood that the requirement will strain your system resources—for example, by demanding more tape units than are available—consider regularly merging multiple image copies into one copy.

Even if you do not periodically merge multiple image copies into one copy when there are not enough tape units, RECOVER TABLESPACE can still attempt to recover the object. RECOVER dynamically allocates the full image copy and attempts to allocate dynamically all the incremental image copy data sets. If every incremental copy can be allocated, recovery proceeds to merge pages to table spaces and apply the log. If a point is reached where an incremental copy cannot be allocated, the log RBA of the last successfully allocated data set is noted.

Attempts to allocate incremental copies cease, and the merge proceeds using only the allocated data sets. The log is applied from the noted RBA, and the incremental image copies that were not allocated are simply ignored.

After running LOAD or REORG: Primary and secondary image copies are recommended after a LOAD or REORG operation specified with LOG NO when an inline copy is not created, so if the primary image copy is not available, fallback recovery using the secondary image copy is possible.

Creating a point of recovery: If you use COPY SHRLEVEL REFERENCE to copy a list of objects which contains all referentially related structures, you do not need to QUIESCE these objects first in order to create a consistent point of recovery.

You should quiesce and copy both the base table space and the LOB table space at the same time to establish a recoverable point of consistency. Be aware that QUIESCE does not create a recoverable point for a LOB table space that contain LOBs defined with LOG NO.

Setting and clearing the informational COPY-pending status: The following utilities can place an index that was defined with the COPY YES attribute in the informational COPY-pending (ICOPY) status:

- REORG INDEX
- REORG TABLESPACE LOG YES or NO
- LOAD TABLE LOG YES or NO
- REBUILD INDEX

After the utility processing completes, take a full image copy of the index space so that the index space is recoverable using the RECOVER utility. If you need to recover an index that did not have a full image copy taken, use the REBUILD INDEX utility to rebuild data from the table space.

Improving performance

A full image copy and subsequent incremental image copies can be merged into a new full copy by running MERGECOPY. After reorganizing a table space, the first image copy *must* be a full image copy.

The decision whether to run a full or an incremental image copy must not be based on the number of rows updated since the last image copy was taken. Instead, it must be based on the percentage of pages containing at least one updated record (not the number of records updated). Regardless of the size of the table, if more than 50% of the pages contain updated records, use full image copy (this saves the cost of a subsequent MERGECOPY). To find the percentage of changed pages, you can execute COPY with the CHANGELIMIT REPORTONLY option. Alternatively, you can execute COPY CHANGELIMIT to allow COPY to determine whether a full or incremental copy is required; see “Specifying conditional image copies” on page 98 for more information.

Using data compression can improve COPY performance because COPY does not decompress data. The performance improvement is proportional to the amount of compression.

Considerations for running COPY

This section describes additional points to keep in mind when running COPY.

Copying table spaces with mixed volume IDs

You cannot copy a table space or index space that uses a storage group that is defined with mixed specific and non-specific volume IDs using CREATE

STOGROUP or ALTER STOGROUP. If you specify such a table space or index space, the job terminates and you receive error message DSNU419I.

Defining generation data groups

We recommend using generation data groups to hold image copies, because their use automates the allocation of data set names and the deletion of the oldest data set. When you define the generation data group:

- You can specify that the oldest data set is automatically deleted when the maximum number of data sets is reached. If you do that, make the maximum number large enough to cover all recovery requirements. When data sets are deleted, use the MODIFY utility to delete the corresponding rows in SYSIBM.SYSCOPY.
- Make the limit number of generation data sets equal to the number of copies to keep. Use NOEMPTY to avoid deleting all the data sets from the integrated catalog facility catalog when the limit is reached.

Attention: Do not take incremental image copies when using generation data groups unless data pages have changed. When you use generation data groups, taking an incremental image copy when no data pages have changed results in the following:

- The new image copy data set is empty
- No SYSCOPY record is inserted for the new image copy data set
- Your oldest image copy is deleted

See “Using conditional copy with generation data groups (GDGs)” on page 99 for guidance on executing COPY with the CHANGELIMIT and REPORTONLY options to ensure that you do not create empty image copy data sets when using GDGs.

Using DB2 with DFSMS products

If image copy data sets are managed by HSM or SMS, all data sets are cataloged.

If you plan to use SMS, catalog all image copies. Never maintain cataloged and un-cataloged image copies with the same name.

Putting image copies on tape

Do not combine a full image copy and incremental image copies for the same table space on one tape volume. If you do, the RECOVER TABLESPACE utility cannot allocate the incremental image copies.

Copying a LOB table space

Both full and incremental image copies are supported for a LOB table space, as well as SHRLEVEL REFERENCE, SHRLEVEL CHANGE, and the CONCURRENT options. COPY without the CONCURRENT option does not copy empty or unformatted data pages for a LOB table space.

Terminating or restarting COPY

For instructions on restarting a utility job, see “Restarting an online utility” on page 46.

Warning against TERM UTILITY

We do *not* recommend stopping a COPY job with the TERM UTILITY command. If you issue TERM UTILITY while COPY is in the active or stopped state, DB2 inserts an ICTYPE='T' record in the SYSIBM.SYSCOPY catalog table for each object COPY had started processing, but not yet completed. For copies made with SHRLEVEL REFERENCE, it is possible that some objects in the list might not have

a 'T' record. For SHRLEVEL CHANGE, some objects might have a valid 'F', 'I', or 'T' record, or no record at all. The COPY utility does not allow you to take an incremental image copy if a 'T' record exists. To reset the status, you must make a full image copy.

Use restart current instead, because it:
- Is valid for full image copies and incremental copies
- Is valid for a single job step with several COPY statements
- Is valid for a list of objects
- Requires a minimum of re-processing
- Keeps the DB2 catalog and the integrated catalog facility catalog in agreement

DB2 uses the same image copy data set when you RESTART from the last commit point. Therefore, specify DISP=(MOD,CATLG,CATLG) on your DD statements. You cannot use RESTART(PHASE) for any COPY job.

Implications of DISP on the DD statement

If you terminate a COPY job that uses the parameter DISP=(MOD,CATLG,CATLG), then:

- If there is only one COPY statement, no row is written to SYSIBM.SYSCOPY, but an image copy data set has been created and is cataloged in the integrated catalog facility catalog. You should delete that data set.
- If there are several COPY statements in one COPY job step, a row for each successfully completed copy is written into SYSIBM.SYSCOPY. However, all the image copy data sets have been created and cataloged. You should delete all image copy data sets not recorded in SYSIBM.SYSCOPY.

For considerations on the disposition setting when restartng a COPY utility job that
uses DFSMS concurrent copy, see “Using DFSMS concurrent copy” on page 97.

Restarting with a new data set

If you define a new output data set for a current restart, complete the following actions before restarting the COPY job:

1. Copy the failed copy output to the new data set.
2. Delete the old data set.
3. Rename the new data set to use the old data set name.

Restarting a COPY job

You cannot use RESTART(PHASE) for any COPY job. If you do *not* use the -TERM UTILITY command, you can use RESTART to restart the job from the last commit point of the current table space or index space.

Restarting COPY after an out of space condition

See “Restarting after the output data set is full” on page 47 for guidance in restarting COPY from the last commit point after receiving an out of space condition.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Claims and drains

Table 13 shows which claim classes COPY claims and drains and any restrictive status the utility sets on the target table space.

Table 13. Claim classes of COPY operations. Use of claims and drains; restrictive states set on the target object.

Target	SHRLEVEL REFERENCE	SHRLEVEL CHANGE
Table space, Index space or partition	DW UTRO	CR UTRW[1]

Legend:
- DW - Drain the write claim class - concurrent access for SQL readers
- CR - Claim the read claim class
- UTRO - Utility restrictive state - read only access allowed
- UTRW - Utility restrictive state - read/write access allowed

Notes:

1. If the target object is a segmented table space, SHRLEVEL CHANGE does not allow you to concurrently execute an SQL *searched* DELETE without the WHERE clause.

COPY does not set a utility restrictive state if the target object is DSNDB01.SYSUTILX.

Compatibility

Table 14 documents which utilities can run concurrently with COPY on the same target object. The target object can be a table space, an index space, or a partition of a table space or index space. If compatibility depends on particular options of a utility, that is also shown.

Table 14. COPY compatibility

Action	COPY INDEXSPACE SHRLEVEL REFERENCE	COPY INDEXSPACE SHRLEVEL CHANGE	COPY TABLESPACE SHRLEVEL REFERENCE	COPY TABLESPACE SHRLEVEL CHANGE
CHECK DATA	Yes	Yes	No	No
CHECK INDEX	Yes	Yes	Yes	Yes
CHECK LOB	Yes	Yes	Yes	Yes
COPY INDEXSPACE	No	No	Yes	Yes
COPY TABLESPACE	Yes	Yes	No	No
DIAGNOSE	Yes	Yes	Yes	Yes
LOAD	No	No	No	No
MERGECOPY	No	No	No	No
MODIFY	No	No	No	No
QUIESCE	Yes	No	Yes	No
REBUILD INDEX	No	No	Yes	Yes
RECOVER INDEX	No	No	Yes	Yes
RECOVER TABLESPACE	Yes	Yes	No	No
REORG INDEX	No	No	Yes	Yes

Table 14. COPY compatibility (continued)

Action	COPY INDEXSPACE SHRLEVEL REFERENCE	COPY INDEXSPACE SHRLEVEL CHANGE	COPY TABLESPACE SHRLEVEL REFERENCE	COPY TABLESPACE SHRLEVEL CHANGE
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No	No	No	No
REORG TABLESPACE UNLOAD ONLY or EXTERNAL	Yes	Yes	Yes	Yes
REPAIR LOCATE by KEY, RID, or PAGE DUMP or VERIFY	Yes	Yes	Yes	Yes
REPAIR LOCATE by KEY or RID DELETE or REPLACE	No	No	No	No
REPAIR LOCATE TABLESPACE PAGE REPLACE	Yes	Yes	No	No
REPAIR LOCATE INDEX PAGE REPLACE	No	No	Yes	No
REPORT	Yes	Yes	Yes	Yes
RUNSTATS INDEX	Yes	Yes	Yes	Yes
RUNSTATS TABLESPACE	Yes	Yes	Yes	Yes
STOSPACE	Yes	Yes	Yes	Yes
UNLOAD	Yes	Yes	Yes	Yes

To run on DSNDB01.SYSUTILX, COPY must be the only utility in the job step. Further, if SHRLEVEL REFERENCE is specified, the COPY job of DSNDB01.SYSUTILX must be the only utility running in the Sysplex.

COPY on SYSUTILX is an “exclusive” job; such a job can interrupt another job between job steps, possibly causing the interrupted job to time out.

Sample control statements

In some cases, a COPY utility job might be run more than once. To facilitate avoiding duplicate image copy data sets, a DSN qualifier is used in the following examples. See the description of the COPYDDN parameter in “Option descriptions” on page 84 for further information.

Example 1: Full image copy. Make a full image copy of table space DSN8S71E in database DSN8D71A.

```
//STEP1    EXEC DSNUPROC,UID='IUJMU111.COPYTS',
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//SYSIN    DD *
//SYSCOPY  DD DSN=COPY001F.IFDY01,UNIT=SYSDA,VOL=SER=CPY01I,
//         SPACE=(CYL,(15,1)),DISP=(NEW,CATLG,CATLG)
//SYSIN DD *
COPY TABLESPACE DSN8D71A.DSN8S71E
```

Example 2: Full image copies of a list of objects. Make full image copies at the local site and recovery site, along with backups, of the following objects:

- table space DSN8D71A.DSN8S71D, and its indexes:
  - DSN8710.XDEPT1
  - DSN8710.XDEPT2
  - DSN8710.XDEPT3
- table space DSN8D71A.DSN8S71E, and its indexes

Do not allow updates during the copy, and process four objects in parallel. As the copy of an object completes, the next object in the list begins processing in parallel until all the objects have been processed.

This COPY job creates a point of consistency for the table spaces and their indexes. You can subsequently use the RECOVER utility with the TOCOPY option to recover all of these objects; see page 302 for an example.

```
COPY
   TABLESPACE DSN8D71A.DSN8S71D
        COPYDDN (COPY1,COPY2)
        RECOVERYDDN (COPY3,COPY4)
  INDEX DSN8710.XDEPT1
        COPYDDN (COPY5,COPY6)
        RECOVERYDDN (COPY7,COPY8)
  INDEX DSN8710.XDEPT2
        COPYDDN (COPY9,COPY10)
        RECOVERYDDN (COPY11,COPY12)
  INDEX DSN8710.XDEPT3
        COPYDDN (COPY13,COPY14)
        RECOVERYDDN (COPY15,COPY16)
  TABLESPACE DSN8D71A.DSN8S71E
        COPYDDN (COPY17,COPY18)
        RECOVERYDDN (COPY19,COPY20)
  INDEX DSN8710.XEMP1
        COPYDDN (COPY21,COPY22)
        RECOVERYDDN (COPY23,COPY24)
  INDEX DSN8710.XEMP2
        COPYDDN (COPY25,COPY26)
        RECOVERYDDN (COPY27,COPY28)
 PARALLEL(4)
 SHRLEVEL REFERENCE
```

Example 3: Copies for local site and recovery site. Make full image copies of table space DSN8S71C in database DSN8D71P at the local site and the recovery site.

```
//STEP1    EXEC DSNUPROC,UID='IUJMU111.COPYLST',
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//COPY1    DD DSN=IUJMU111.COPYLST.STEP1.COPY1,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//COPY2    DD DSN=IUJMU111.COPYLST.STEP1.COPY2,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(2000,(20,20),,,ROUND)
//COPY3    DD DSN=IUJMU111.COPYLST.STEP1.COPY3,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(2000,(20,20),,,ROUND)
//COPY4    DD DSN=IUJMU111.COPYLST.STEP1.COPY4,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(2000,(20,20),,,ROUND)
//SYSIN    DD *
```

```
COPY TABLESPACE DSN8D71P.DSN8S71C
       COPYDDN     (COPY1,COPY2)
       RECOVERYDDN (COPY3,COPY4)
```

Example 4: Incremental copy with updates allowed. Make incremental image copies of table space DSN8S71D in database DSN8D71A, allowing update activity to occur during the copy process.

```
//STEP1    EXEC DSNUPROC,UID='IUJMU111.COPYLSTI',
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//SYSCOPY  DD DSN=IUJMU111.COPYLSTI.STEP1.CPY01I,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//COPY2    DD DSN=IUJMU111.COPYLSTI.STEP1.CPY02I,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(2000,(20,20),,,ROUND)
//COPY3    DD DSN=IUJMU111.COPYLSTI.STEP1.CPY03I,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(2000,(20,20),,,ROUND)
//COPY4    DD DSN=IUJMU111.COPYLSTI.STEP1.CPY04I,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(2000,(20,20),,,ROUND)
//SYSIN    DD *
COPY TABLESPACE DSN8D71A.DSN8S71D
     COPYDDN      (SYSCOPY,COPY2)
     RECOVERYDDN  (COPY3,COPY4)
     FULL NO
     SHRLEVEL CHANGE
```

Example 5: Invoking DFSMS concurrent copy with the COPY utility. Copy a table space, using the CONCURRENT option to execute DFSMS concurrent copy. Use a DSSPRINT DD card for message output. Specify a specific volume, DB2CC5, for the output new image copy.

```
//COPY     EXEC DSNUPROC,SYSTEM=DSN
//SYSCOPY1 DD DSN=COPY1,DISP=(NEW,CATLG,CATLG),
//         SPACE=(4000,(20,20),,,ROUND),UNIT=SYSDA,VOL=SER=DB2CC5
//SYSPRINT DD DSN=COPY1.PRINT1,DISP=(NEW,CATLG,CATLG),
//         SPACE=(4000,(20,20),,,ROUND),UNIT=SYSDA,VOL=SER=DB2CC5
//DSSPRINT DD DSN=COPY1.PRINT2,DISP=(NEW,CATLG,CATLG),
//         SPACE=(4000,(20,20),,,ROUND),UNIT=SYSDA,VOL=SER=DB2CC5
//SYSIN    DD *
COPY TABLESPACE DBASE1AA.TABLESPC
  COPYDDN     (SYSCOPY1)
  CONCURRENT
```

Example 6: Invoking DFSMS concurrent copy with the COPY utility using FILTER. Copy a list of table spaces, using the CONCURRENT and FILTERDDN options to create a single ″DUMP″ statement for DFSMS concurrent copy, allowing maximum availability.

```
//SYSCOPY  DD DSN=CONCOPY.WFILT,DISP=(MOD,CATLG,CATLG),
//            UNIT=SYSDA,SPACE=(CYL,(42,5),RLSE)
//FILT     DD DSN=FILT.TEST1,DISP=(MOD,CATLG,DELETE),
//            UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE)
//SYSIN    DD *
COPY TABLESPACE TS1
     TABLESPACE TS2
     TABLESPACE TS3
FILTERDDN(FILT)
COPYDDN(SYSCOPY)
CONCURRENT
    SHRLEVEL REFERENCE
```

Example 7: Invoking DFSMS concurrent copy with a list. Copy a list of table spaces, using the CONCURRENT option to execute DFSMS concurrent copy. Allow update activity during the COPY operation.

```
//STEP1    EXEC DSNUPROC,UID='IUJMU111.COPYLST',
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//COPY1    DD DSN=IUJMU111.COPYLST.STEP1.TS1,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//COPY2    DD DSN=IUJMU111.COPYLST.STEP1.TS2,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(2000,(20,20),,,ROUND)
//COPY3    DD DSN=IUJMU111.COPYLST.STEP1.TS3,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(2000,(20,20),,,ROUND)
//SYSIN    DD *
COPY TABLESPACE DBAU2901.TPAU2901
     COPYDDN(COPY1)
     TABLESPACE DBAU2901.TLAU2902
     COPYDDN(COPY2)
     TABLESPACE DBAU2901.TSAU2903
     COPYDDN(COPY3)
     CONCURRENT SHRLEVEL CHANGE
```

Example 8: Report image copy information for a table space. Recommend a full image copy if the percent of changed pages is equal to or greater than 40 percent. Recommend an incremental image copy if the percent of changed pages is greater than 10 and less than 40 percent. Recommend no image copy if the percent of changed pages is 10 percent or less.

```
COPY TABLESPACE DSN8D71P.DSN8S71C CHANGELIMIT(10,40) REPORTONLY
```

Example 9: Make a conditional image copy. Take a full image copy of a table space if the number of changed pages is equal to or greater than 5 percent. Take an incremental image copy if the percent of changed pages is greater than 0 and less than 5 percent. If no pages have changed, do not take an image copy.

```
COPY TABLESPACE DSN8D71P.DSN8S71C CHANGELIMIT(5)
```

Example 10: Copying LOB table spaces together with related objects. Take a full image copy of base table space TPIQUD01 and LOB table spaces TLIQUDA1, TLIQUDA2, TLIQUDA3, and TLIQUDA4 in database DBIQUD01 if the number of changed pages is equal to or greater than the decimal percentage values specified for each object. Take an incremental image copy if the percent of changed pages falls in the range between the specified decimal percentage values. If no pages have changed, do not take an image copy. Also take full image copies of index spaces IPIQUD01, IXIQUD02, IUIQUD03, IXIQUDA1, IXIQUDA2, IXIQUDA3, and IXIQUDA4.

```
COPY
     TABLESPACE DBIQUD01.TPIQUD01 DSNUM ALL CHANGELIMIT(3.3,6.7)
       COPYDDN(COPYTB1)
     TABLESPACE DBIQUD01.TLIQUDA1 DSNUM ALL CHANGELIMIT(7.9,25.3)
       COPYDDN(COPYTA1)
     TABLESPACE DBIQUD01.TLIQUDA2 DSNUM ALL CHANGELIMIT(2.2,4.3)
       COPYDDN(COPYTA2)
     TABLESPACE DBIQUD01.TLIQUDA3 DSNUM ALL CHANGELIMIT(1.2,9.3)
       COPYDDN(COPYTA3)
     TABLESPACE DBIQUD01.TLIQUDA4 DSNUM ALL CHANGELIMIT(2.2,4.0)
       COPYDDN(COPYTA4)
     INDEXSPACE DBIQUD01.IPIQUD01 DSNUM ALL
       COPYDDN(COPYIX1)
     INDEXSPACE DBIQUD01.IXIQUD02 DSNUM ALL
       COPYDDN(COPYIX2)
```

```
    INDEXSPACE DBIQUD01.IUIQUD03 DSNUM ALL
      COPYDDN(COPYIX3)
    INDEXSPACE DBIQUD01.IXIQUDA1 DSNUM ALL
      COPYDDN(COPYIXA1)
    INDEXSPACE DBIQUD01.IXIQUDA2 DSNUM ALL
      COPYDDN(COPYIXA2)
    INDEXSPACE DBIQUD01.IXIQUDA3 DSNUM ALL
      COPYDDN(COPYIXA3)
    INDEXSPACE DBIQUD01.IXIQUDA4 DSNUM ALL
      COPYDDN(COPYIXA4)
SHRLEVEL REFERENCE
```

Chapter 10. COPYTOCOPY

The COPYTOCOPY utility makes image copies from an image copy that was taken by the COPY utility. This includes inline copies made by REORG or LOAD utility. Starting with either the local primary or recovery site primary copy, COPYTOCOPY can make up to three copies of one or more of the following types of copies:

- local primary
- local backup
- recovery site primary
- recovery site backup

The copies are used by the RECOVER utility when recovering a table space or index space to the most recent time or to a previous time. These copies can also be used by MERGECOPY, UNLOAD, and possibly a subsequent COPYTOCOPY execution.

For a diagram of COPYTOCOPY syntax and a description of available options, see "Syntax and options of the control statement" on page 110. For detailed guidance on running this utility, see "Instructions for running COPYTOCOPY" on page 115.

Output: Output from the COPYTOCOPY utility consists of:

- Up to three sequential data sets containing the image copy.
- Rows in the SYSIBM.SYSCOPY catalog table that describe the image copy data sets available to the RECOVER utility. It is your installation's responsibility to ensure that these data sets are available if the RECOVER utility requests them.

The entries for SYSIBM.SYSCOPY columns remain the same as the original entries in the SYSIBM.SYSCOPY row when the COPY utility recorded them. The COPYTOCOPY job inserts values in the columns DSNAME, GROUP_MEMBER, JOBNAME, AUTHID, DSVOLSER and DEVTYPE.

Restrictions: COPYTOCOPY does not support the following catalog and directory objects:

- DSNDB01.SYSUTILX, and its indexes
- DSNDB01.DBD01, and its indexes
- DSNDB06.SYSCOPY, and its indexes

An image copy in which the CONCURRENT option of COPY was specified is not supported as an input to COPYTOCOPY.

COPYTOCOPY does not check recoverability of an object.

Related information: See Part 4 (Volume 1) of *DB2 Administration Guide* for uses of COPYTOCOPY in the context of planning for database recovery.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:

- IMAGCOPY privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority

© Copyright IBM Corp. 1983, 2001

An ID with installation SYSOPR authority can also execute COPYTOCOPY, but only on a table space in the DSNDB01 or DSNDB06 database.

Execution phases of COPYTOCOPY: The COPYTOCOPY utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup
CPY2CPY	Copying of an image copy
UTILTERM	Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

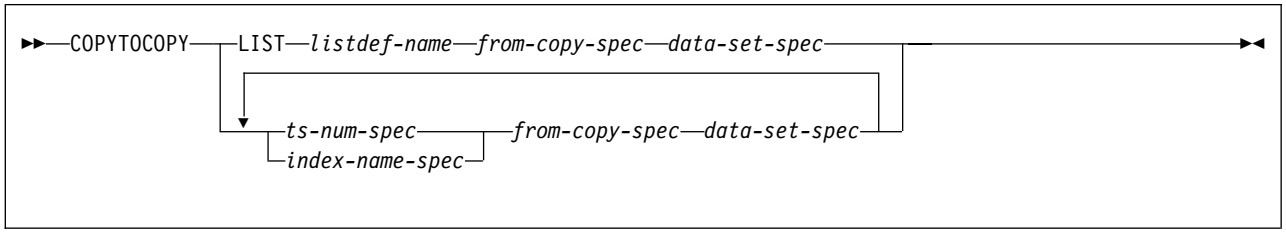

ts-num-spec:

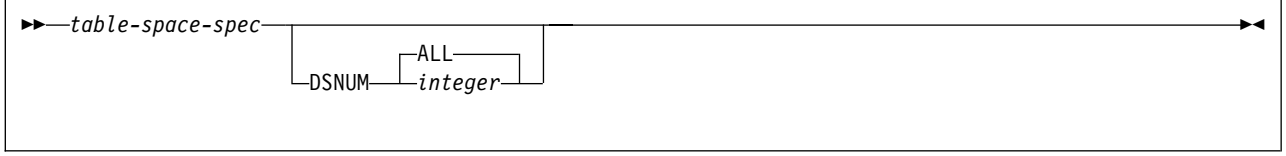

table-space-spec:

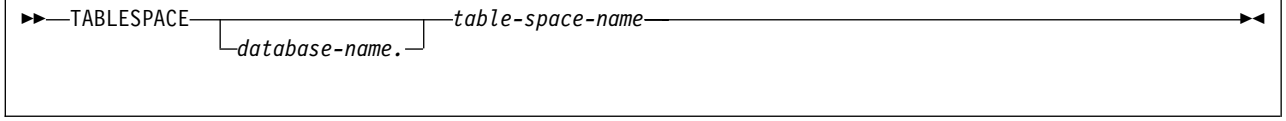

index-name-spec:

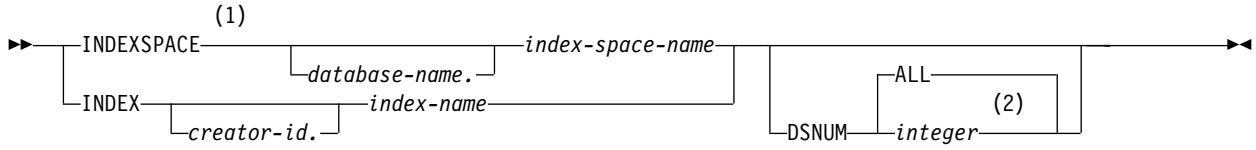

Notes:

1 INDEXSPACE is the preferred specification.

2 Not valid for nonpartitioning indexes.

from-copy-spec:

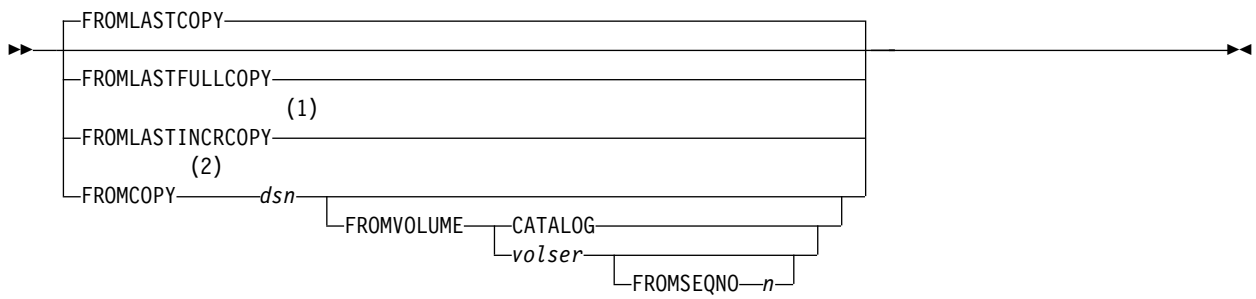

Notes:

1 Not valid with the INDEXSPACE or INDEX keyword.

2 Not valid with the LIST keyword.

data-set-spec:

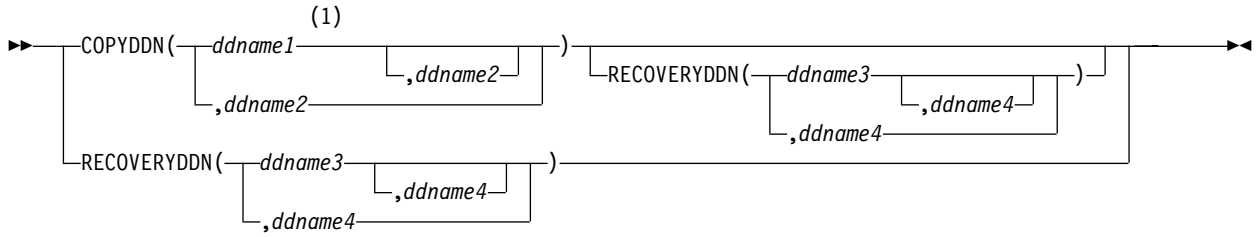

Notes:

1 Use this option if you want to make a local site primary copy from one of the recovery site copies.

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each COPYTOCOPY control statement. Do not specify LIST with either the INDEX keywords or the TABLESPACE. For more information about LISTDEF specifications, see "Generating object lists using LISTDEF" on page 144.

TABLESPACE
Specifies the table space (and, optionally, the database it belongs to) that is to be copied.

database-name
Is the name of the database the table space belongs to.

The **default** is **DSNDB04**.

table-space-name
Is the name of the table space to be copied.

INDEXSPACE *database-name.index-space-name*
Specifies the qualified name of the index space that is to be copied; the name is obtained from the SYSIBM.SYSINDEXES table. The index space specified must be defined with the COPY YES attribute.

database-name
Optionally specifies the name of the database the index space belongs to.

The **default** is **DSNDB04**.

index-space-name
Specifies the name of the index space to be copied.

INDEX *creator-id.index-name*
Specifies the index to be copied.

creator-id
Optionally specifies the creator of the index.

The **default** is the user identifier for the utility.

index-name
Specifies the name of the index that is to be copied.

DSNUM Identifies a partition or data set, within the table space or the index space, that is to be copied. The keyword ALL specifies that the entire table space or index space is copied.

ALL Copies the entire table space or index space.

The **default** is **ALL**.

You must use ALL for a nonpartitioning index.

integer
Is the number of a partition or data set to be copied.

An integer value is not valid for nonpartitioning indexes.

For a partitioned table space or index space, the integer is its partition number. The maximum is 254.

For a nonpartitioned table space, find the integer at the end of the data set name as cataloged in the VSAM catalog. The data set name has this format:

catname.DSNDB*x*.*dbname*.*spacename*.*y*0001.A*nnn*

where:

catname	The VSAM catalog name or alias
x	C or D
dbname	The database name
spacename	The table space or index space name
y	I or J
nnn	The data set integer.

Specifying DSNUM(ALL) or defaulting to it will cause COPYTOCOPY to look for an input image copy taken at the entire table space or index space level.

FROMLASTCOPY

Specifies the most recent image copy that was taken for the table space or index space to be the input to the COPYTOCOPY utility. This could be a full image copy or incremental copy retrieved from SYSIBM.SYSCOPY.

FROMLASTFULLCOPY

Specifies the most recent full image copy that was taken for the object to be the input to COPYTOCOPY job.

FROMLASTINCRCOPY

Specifies the most recent incremental image copy that was taken for the object to be the input to COPYTOCOPY job.

FROMLASTINCRCOPY is not valid with the INDEXSPACE or INDEX keyword. If FROMLASTINCRCOPY is specified for an INDEXSPACE or INDEX, COPYTOCOPY will use the last full copy taken, if one is available.

FROMCOPY *dsn*

Specifies a particular image copy data set (*dsn*) as the input to COPYTOCOPY job. This option is not valid for LIST.

If the image copy data set is an MVS generation data set, then supply a fully qualified data set name including the absolute generation and version number. If the image copy data set is not a generation data set and there is more than one image copy data set with the same data set name, use the FROMVOLUME option to identify the data set exactly.

FROMVOLUME

Identifies the image copy data set.

CATALOG

Identifies the data set as cataloged. Use this option only for an image copy that was created as a cataloged data set. (Its volume serial is not recorded in SYSIBM.SYSCOPY.)

COPYTOCOPY refers to the SYSIBM.SYSCOPY catalog table during execution. If you use FROMVOLUME CATALOG, the data set must be cataloged. If you remove the data set from the catalog after creating it, you must catalog the data set again to make it consistent with the record for this copy that appears in SYSIBM.SYSCOPY.

vol-ser
Identifies the data set by an alphanumeric volume serial identifier of its first volume. Use this option only for an image copy that was created as a noncataloged data set. Specify the first *vol-ser* in the SYSCOPY record to locate a data set stored on multiple tape volumes.

FROMSEQNO *n*
Identifies the image copy data set by its file sequence number. *n* is the file sequence number.

COPYDDN *ddname1,ddname2*
Specifies a DD name or a TEMPLATE name for the primary (*ddname1*) and backup (*ddname2*) copied data sets for the image copy at the local site. If *ddname2* is specified by itself, COPYTOCOPY expects local site primary to exist. If it does not exist, error message DSNU1401 will be issued and process for the object will be terminated.

ddname is the DD name.

Recommendation: Catalog all of your image copy data sets.

You cannot have duplicate image copy data sets. If the DD statement identifies a noncataloged data set with the same name, volume serial, and file sequence number as one already recorded in SYSIBM.SYSCOPY, a message is issued and no copy is made. If it identifies a cataloged data set with only the same name, no copy is made. For cataloged image copy data sets, CATLG must be specified for the normal termination disposition in the DD statement; for example, DISP=(MOD,CATLG,CATLG). The DSVOLSER field of the SYSIBM.SYSCOPY entry will be blank.

When the image copy data set is going to a tape volume, the VOL=SER parameter should be specified on the DD statement.

The COPYDDN keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

RECOVERYDDN *ddname3,ddname4*
Specifies a DD name or a TEMPLATE name for the primary (*ddname3*) and backup (*ddname4*) copied data sets for the image copy at the recovery site. If *ddname4* is specified by itself, COPYTOCOPY expects recovery site primary to exist. If it does not exist, error message DSNU1401 will be issued and process for the object will be terminated.

ddname is the DD name.

You cannot have duplicate image copy data sets. The same rules apply for RECOVERYDDN as for COPYDDN.

The RECOVERYDDN keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

Instructions for running COPYTOCOPY

To run COPYTOCOPY, you must:

1. Read "Before running COPYTOCOPY" in this chapter.
2. Prepare the necessary data sets, as described in "Data sets used by COPYTOCOPY".
3. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL and control statement for COPYTOCOPY, see "Sample control statement" on page 120.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 116. (For a complete description of the syntax and options for COPYTOCOPY, see "Syntax and options of the control statement" on page 110.)
5. Plan for restart if the COPYTOCOPY job does not complete, as described in "Terminating or restarting COPYTOCOPY" on page 119.
6. Run COPYTOCOPY.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Before running COPYTOCOPY

Check the compatibility table in "Concurrency and compatibility" on page 119 if you want to run other jobs concurrently on the same target objects.

Data sets used by COPYTOCOPY

Table 15 describes the data sets required for COPYTOCOPY. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 15. Data sets used by COPYTOCOPY

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
Output Copies	From one to three output data sets to contain the resulting image copy data sets. Their DD names are specified with the COPYDDN and RECOVERYDDN options of the utility control statement.	Yes

The following objects are named in the utility control statement and do not require DD cards in the JCL:

Table space or Index space
The object to be copied is named in the COPYTOCOPY control statement and is accessed through the DB2 catalog. (If you want to copy only certain partition in a partitioned table space, you must use the DSNUM option in the control statement.)

DB2 catalog objects
COPYTOCOPY records each copy it makes in the DB2 catalog table SYSIBM.SYSCOPY.

Input image copy data set
This information is accessed through the DB2 catalog. However, if you would like to preallocate your image copy data sets by using DD cards, refer to "Retaining tape mounts" on page 118 for more information.

Output data set size: Image copies are written to sequential non-VSAM data sets. To find the approximate size of the image copy data set for a table space in bytes, use the following procedure:

1. Find the *high allocated page number* from the COPYPAGESF column of SYSIBM.SYSCOPY or from information in the VSAM catalog data set.
2. Multiply the high allocated page number by the page size.

Another option is to look at the size of the input image copy.

JCL parameters: You can specify a block size for the output by using the BLKSIZE parameter on the DD statement for the output data set. Valid block sizes are multiples of 4096 bytes.

Cataloging image copies: To catalog your image copy data sets, use the DISP=(NEW,CATLG,CATLG) parameter in the DD statement named by the COPYDDN or RECOVERYDDN option. After the image copy is taken, the DSVOLSER column of the row inserted into SYSIBM.SYSCOPY contains blanks.

Duplicate image copy data sets are not allowed. If there is a cataloged data set already recorded in SYSCOPY with the same name as the new image copy data set, a message is issued and the copy is not made.

When RECOVER locates the entry in SYSCOPY, it uses the MVS catalog to allocate the required data set. *If you have uncataloged the data set, the allocation fails.* In that case, the recovery can still go forward; RECOVER searches for a previous image copy. But even if it finds one, it must use correspondingly more of the log to recover. It is to your benefit, and it is your responsibility, to keep the MVS catalog consistent with SYSIBM.SYSCOPY about existing image copy data sets.

Creating the control statement

See "Syntax and options of the control statement" on page 110 for COPYTOCOPY syntax and option descriptions. See "Sample control statement" on page 120 for examples of COPYTOCOPY usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

"Copying full or incremental image copies" on page 117
"Copying incremental image copies" on page 117
"Using more than one COPYTOCOPY statement" on page 117

"Copying inline copy made by REORG with part range"
"Copying from a specific image copy"
"Using TEMPLATE with COPYTOCOPY" on page 118

Copying full or incremental image copies

You can copy a full image copy or an incremental image copy using FROMLASTCOPY keyword. If you do not specify FROMLASTCOPY, it will be used by default, as shown below. COPYTOCOPY makes a backup copy of a most recent full image copy or an incremental image copy of the table space DSN8S71E in database DSN8D71A:

```
COPYTOCOPY TABLESPACE DSN8D71A.DSN8S71E
     COPYDDN(,DDNNAME2)
```

The COPYTOCOPY utility makes a copy from an existing image copy and writes pages from the image copy to the output data sets. The JCL for the utility job must include DD statements or a template for the output data sets. If the object consists of multiple data sets and all are copied in one run, the copies reside in one physical sequential output data set.

Copying incremental image copies

An incremental image copy is a copy of the pages that have been changed since the last full or incremental image copy. To make a copy of an incremental image copy, use the keyword FROMLASTINCRCOPY.

The following example make a local site backup, recovery site primary and recovery site backup from an incremental image copy.

```
COPYTOCOPY TABLESPACE DSN8D71A.DSN8S71E
  FROMLASTINCRCOPY
  COPYDDN(,COPY2)
  RECOVERYDDN(COPY3,COPY4)
```

Using more than one COPYTOCOPY statement

You can use more than one control statement for COPYTOCOPY in one DB2 utility job step. After each COPYTOCOPY statement has executed successfully:

- A row referring to the image copy is recorded in SYSIBM.SYSCOPY table.
- The image copy data set is valid and available for RECOVER, MERGECOPY, COPYTOCOPY and UNLOAD.

If a job step containing more than one COPYTOCOPY statement abends, do not use TERM UTILITY. Restart the job from the last commit point using RESTART instead. Terminating COPYTOCOPY in this case may cause inconsistencies between the ICF catalog and DB2 catalogs if generation data sets are used.

Copying inline copy made by REORG with part range

COPYTOCOPY does not support part range. Specify individual DSNUM(n). From the inline copy, COPYTOCOPY copies only the specified partition into the output image copy data.

Copying from a specific image copy

You can specify a particular image copy to be used as input to COPYTOCOPY using FROMCOPY. The following statement makes three copies of the table space TPA9031C in database DBA90301 from the image copy data set DH109003.COPY1.STEP1.COPY3.

```
COPYTOCOPY TABLESPACE DBA90301.TPA9031C
          FROMCOPY DH109003.COPY1.STEP1.COPY3
          COPYDDN(,COPY2)
          RECOVERYDDN(COPY3,COPY4)
```

If the FROMCOPY keyword is specified and the specified data set is not found in SYSIBM.SYSCOPY, COPYTOCOPY issues DSNU1401I message. Process for the object is then terminated.

Using TEMPLATE with COPYTOCOPY

Template data set name substitution variables will be resolved as usual. COPYTOCOPY does not use the template values of the original COPY utility execution.

Considerations for running COPYTOCOPY

This section describes additional points to keep in mind when running COPYTOCOPY.

Updating SYSIBM.SYSCOPY records

The image copies COPYTOCOPY made are registered in SYSIBM.SYSCOPY to be used by the RECOVER utility. Other utilities can use these copies too. Columns inserted by COPYTOCOPY are the same as those of the original entries in SYSIBM.SYSCOPY row when the COPY utility recorded them. Except for columns GROUP_MEMBER, JOBNAME, AUTHID, DSNAME, DEVTYPE and DSVOLSER, they will be those of the COPYTOCOPY job. When COPYTOCOPY is invoked at the partition level (DSNUM n) and the input data set is an inline copy created by the REORG of a part range, COPYTOCOPY inserts zeros in the HIGHDSNUM and LOWDSNUM columns in the SYSIBM.SYSCOPY record.

Determining the input copy used

If FROMCOPY keyword is not specified, COPYTOCOPY utility uses the following search order to determine the input data set for the utility:

- If you are running the utility at the local site, the search order will be local site primary copy, local site backup copy, recovery site primary copy, recovery site backup copy.
- If you are running the utility at the recovery site, the search order will be recovery site primary copy, recovery site backup copy, local site primary copy, and local site backup copy.

If the input data set cannot be allocated or opened, COPYTOCOPY will attempt to use the next image copy data set, with the same START_RBA value in SYSIBM.SYSCOPY column, in the search order mentioned above.

If FROMCOPY keyword is used, only the specified data set is used as the input to COPYTOCOPY job.

Retaining tape mounts

If the image copy data sets used by COPYTOCOPY are residing on the same tape, you do not need to remove the tape. For noncataloged image copies, specify the following parameters on the DD cards (in this example, the DD cards are *ddname1* and *ddname2*):

```
//ddname1   DD   UNIT=TAPE,DSN=data-set-name1,DISP=(OLD,PASS),LABEL=1,
//               VOL=(,RETAIN,SER=vol-ser)
//ddname2   DD   UNIT=TAPE,DSN=data-set-name2,DISP=(OLD,PASS),LABEL=2,
//               VOL=(,REF=*.ddname1)
```

This example only works for multiple image copies on a single volume. To use multiple image copies on multiple volumes, the image copy data sets must be cataloged. For cataloged image copies on one or more tape volumes, specify the following parameters on the DD cards (in this example, the DD cards are *ddname1*, *ddname2*, and *ddname3*):

```
//ddname1   DD   DSN=data-set-name1,UNIT=3480,DISP=(OLD,PASS),VOL=(,RETAIN),
//                LABEL=(1,SL)
//ddname2   DD   DSN=data-set-name2,UNIT=3480,DISP=(OLD,PASS),VOL=(,RETAIN),
//                LABEL=(2,SL)
//ddname3   DD   DSN=data-set-name3,UNIT=3480,DISP=(OLD,PASS),VOL=(,RETAIN),
//                LABEL=(3,SL)
```

Defining generation data groups

We recommend using generation data groups to hold image copies, because their use automates the allocation of data set names and the deletion of the oldest data set. When you define the generation data group:

- You can specify that the oldest data set is automatically deleted when the maximum number of data sets is reached. If you do that, make the maximum number large enough to cover all recovery requirements. When data sets are deleted, use the MODIFY utility to delete the corresponding rows in SYSIBM.SYSCOPY.
- Make the limit number of generation data sets equal to the number of copies to keep. Use NOEMPTY to avoid deleting all the data sets from the integrated catalog facility catalog when the limit is reached.

Using DB2 with DFSMS products

If image copy data sets are managed by HSM or SMS, all data sets are cataloged.

If you plan to use SMS, catalog all image copies. Never maintain cataloged and uncataloged image copies with the same name.

Putting image copies on tape

Do not combine a full image copy and incremental image copies for the same table space on one tape volume. If you do, the RECOVER TABLESPACE utility cannot allocate the incremental image copies.

Copying a LOB table space

Both full and incremental image copies are supported for a LOB table space.

Terminating or restarting COPYTOCOPY

For instructions on restarting a utility job, see “Restarting an online utility” on page 46.

Restarting a COPYTOCOPY job

You cannot use RESTART(PHASE) for any COPYTOCOPY job. If you do *not* use the -TERM UTILITY command, you can use RESTART(CURRENT) to restart the job. When you use RESTART(CURRENT), COPYTOCOPY will restart from the last commit point with the same image copy data set. Therefore, specify DISP=(MOD,CATLG,CATLG) on your DD statements.

Restarting COPYTOCOPY after an out of space condition

See “Restarting after the output data set is full” on page 47 for guidance in restarting COPYTOCOPY from the last commit point after receiving an out of space condition.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Claims

COPYTOCOPY

Table 16 shows which claim classes COPYTOCOPY claims on the target table space.

Table 16. Claim classes of COPYTOCOPY operations.

Target	COPYTOCOPY
Table space or partition, or index space or partition	UTRW

Legend:
- UTRW - Utility restrictive state - read/write access allowed

Compatibility

Table 17 documents which utilities can run concurrently with COPYTOCOPY on the same target object. The target object can be a table space, an index space, or a partition of a table space or index space. If compatibility depends on particular options of a utility, that is also shown.

Table 17. COPYTOCOPY compatibility

Action	COPYTOCOPY
CHECK DATA	Yes
CHECK INDEX	Yes
CHECK LOB	Yes
COPY	No
DIAGNOSE	Yes
LOAD	No
MERGECOPY	No
MODIFY	No
QUIESCE	Yes
REBUILD INDEX	Yes
RECOVER	No
REORG INDEX	No
REORG TABLESPACE	No
REPAIR	Yes
REPORT	Yes
RUNSTATS INDEX	Yes
RUNSTATS TABLESPACE	Yes
STOSPACE	Yes
UNLOAD	Yes

Sample control statement

Example 1: Make a local backup copy. Make a copy of the most recent full image copy or incremental image copy whichever is recent.

```
//STEP1    EXEC DSNUPROC,UID='DH109001.COPY1',
//         UTPROC='',
//         SYSTEM='V71A'
```

```
//COPY2    DD DSN=DH109001.C2C01.STEP2.COPY2,DISP=(MOD,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(1000,(20,20),,,ROUND)
//SYSIN    DD *

   COPYTOCOPY TABLESPACE DBA90101.TLA9011A COPYDDN(,COPY2)

//
```

Example 2: Use FROMLASTCOPY. Use FROMLASTCOPY to specify the most recent full image copy or incremental image copy to be used as input copy data set. Make local backup, recovery primary and recovery backup copies of table space TPA9012C in database DBA90102.

```
COPYTOCOPY TABLESPACE DBA90102.TPA9012C
                   FROMLASTCOPY COPYDDN(,COPY2)
                   RECOVERYDDN(COPY3,COPY4)
```

Example 3: Use FROMLASTFULLCOPY. Use FROMLASTFULLCOPY to specify the most recent full image copy to be used as input copy data set. Make recovery site primary and recovery site backup copies.

```
COPYTOCOPY TABLESPACE DBA90201.TPA9021C
          FROMLASTFULLCOPY
          RECOVERYDDN(COPY3,COPY4)
```

Example 4: Specify a copy data set for input with FROMCOPY. Make local site backup, recovery site primary and recovery site backup from data set DH109003.COPY1.STEP1.COPY3.

```
COPYTOCOPY TABLESPACE DBA90301.TPA9031C
          FROMCOPY DH109003.COPY1.STEP1.COPY3
          COPYDDN(,COPY2)
          RECOVERYDDN(COPY3,COPY4)
```

Example 5: Identify the image copy data set with FROMVOLUME. Make local site backup from a cataloged data set DH109003.COPY1.STEP1.COPY4.

```
COPYTOCOPY TABLESPACE DBA90302.TLA9032A
          FROMCOPY DH109003.COPY1.STEP1.COPY4
          FROMVOLUME CATALOG
          COPYDDN(,COPY2)
```

Example 6: Specify an uncataloged data set using FROMVOLUME. Make local site backup, recovery site primary and recovery site backup from an uncataloged data set DH109003.COPY1.STEP1.COPY8.

```
//STEP1    EXEC DSNUPROC,UID='DH109003.COPY1',
//         UTPROC='',
//         SYSTEM='SSTR'
//COPY2    DD DSN=DH109003.C2C08.STEP2.COPY2,DISP=(MOD,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(1000,(20,20),,,ROUND)
//COPY3    DD DSN=DH109003.C2C08.STEP2.COPY3,DISP=(MOD,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(1000,(20,20),,,ROUND)
//COPY4    DD DSN=DH109003.C2C08.STEP2.COPY4,DISP=(MOD,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(1000,(20,20),,,ROUND)
//SYSIN    DD *

   COPYTOCOPY TABLESPACE DBA90303.TSA9033B
          FROMCOPY DH109003.COPY1.STEP1.COPY8
          FROMVOLUME SCR03
          COPYDDN(,COPY2)
          RECOVERYDDN(COPY3,COPY4)

//
```

Example 7: Use a list of objects. Make local site backup copies from their image copies.

```
COPYTOCOPY
          TABLESPACE DBA90402.TPA9042C DSNUM 2
              FROMLASTFULLCOPY COPYDDN(,COPY2)
          TABLESPACE DBA90402.TPA9042C DSNUM 3
              FROMLASTCOPY COPYDDN(,COPY3)
          TABLESPACE DBA90402.TPA9042C DSNUM 4
              FROMLASTINCRCOPY COPYDDN(,COPY4)
```

Example 8: Use LISTDEF and TEMPLATE. Copying a list of objects generated with LISTDEF control statement using TEMPLATE. COPYTOCOPY utility makes copies of a LISTDEF list of objects from their most recent image copies. The output data set is named using the format defined by the TEMPLATE control statement.

```
//STEP2    EXEC DSNUPROC,UID='DH109006.COPY1',
//         UTPROC='',
//         SYSTEM='V71A'
//SYSIN    DD *

   OPTIONS PREVIEW
   LISTDEF CPY1 INCLUDE TABLESPACES TABLESPACE DBA906*.T*A906*
           INCLUDE INDEXSPACES COPY YES INDEXSPACE ADMF001.I?A906*
   OPTIONS OFF
   TEMPLATE TMP1 UNIT SYSDA
                 DSN  (DH109006.COPY&LOCREM.&PRIBAC..&SN..T&TIME.)
                 DISP (MOD,CATLG,CATLG)
   COPYTOCOPY LIST CPY1 COPYDDN(TMP1,TMP1)

//
```

Chapter 11. DIAGNOSE

The DIAGNOSE online utility generates information useful in diagnosing problems. It is intended to be used only under the direction of your IBM Support Center.

Interpreting output

One intended use of this utility is to aid in determining and correcting system problems. When diagnosing DB2, you might need to refer to licensed documentation to interpret output from this utility.

For a diagram of DIAGNOSE syntax and a description of available options, see "Syntax and options of the control statement". For detailed guidance on running this utility, see "Instructions for running DIAGNOSE" on page 127.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:

- REPAIR privilege for the database
- DBADM or DBCTRL authority for the database
- SYSCTRL or SYSADM authority.

An ID with installation SYSOPR authority can execute the DIAGNOSE utility on a table space in the DSNDB01 or DSNDB06 database.

An ID with installation SYSADM authority can execute the DIAGNOSE utility with the WAIT statement option.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

© Copyright IBM Corp. 1983, 2001

DIAGNOSE

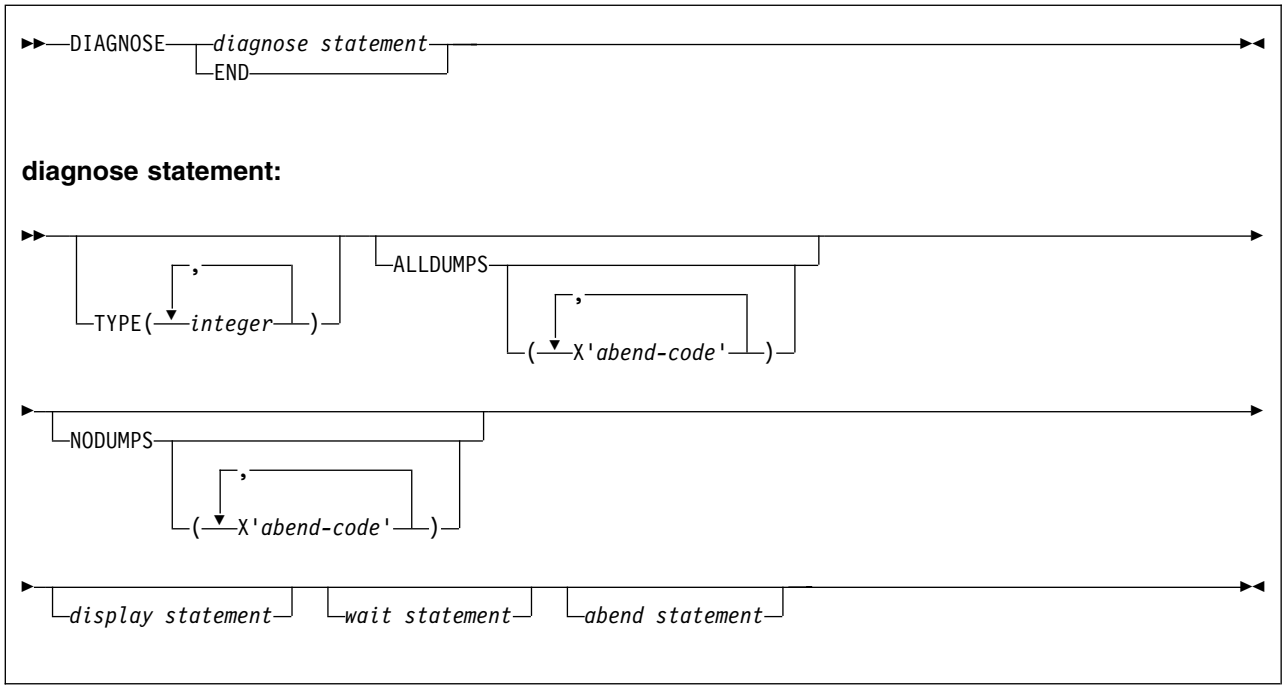
DIAGNOSE
diagnose statement
END
diagnose statement:
TYPE(
,
integer
)
ALLDUMPS
(
,
X'abend-code'
)
NODUMPS
(
,
X'abend-code'
)
display statement
wait statement
abend statement

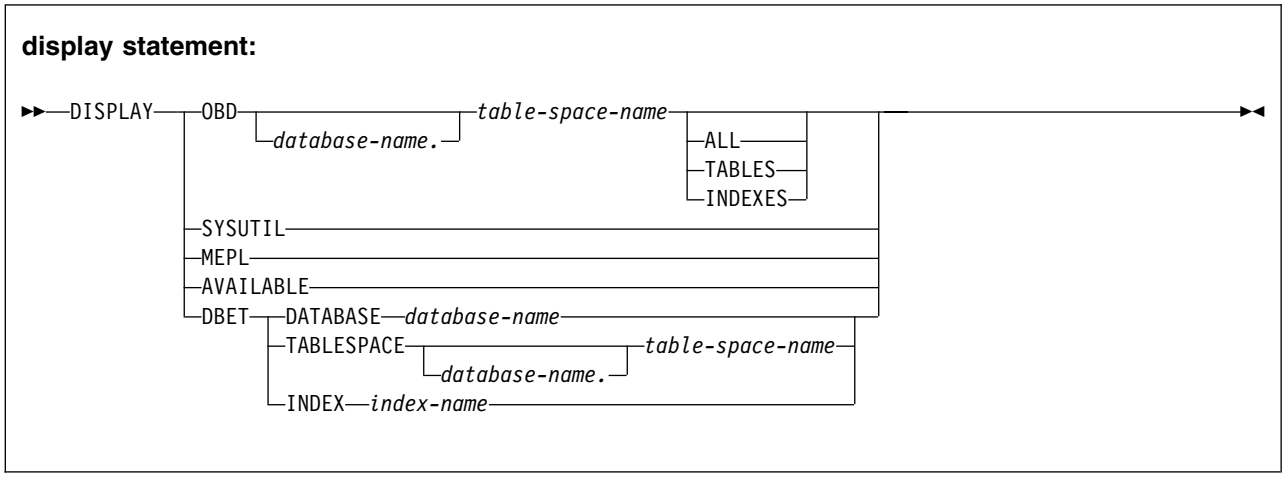
display statement:
DISPLAY
OBD
database-name.
table-space-name
ALL
TABLES
INDEXES
SYSUTIL
MEPL
AVAILABLE
DBET
DATABASE
database-name
TABLESPACE
database-name.
table-space-name
INDEX
index-name

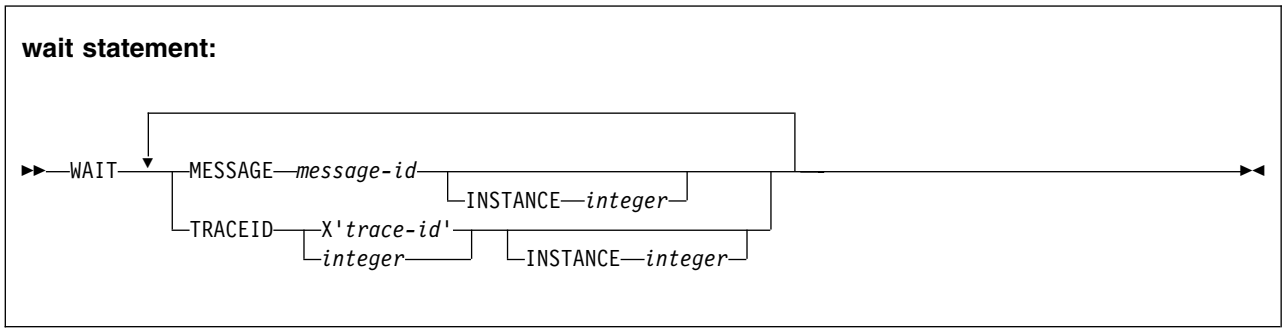
wait statement:
WAIT
MESSAGE
message-id
INSTANCE
integer
TRACEID
X'trace-id'
integer
INSTANCE
integer

abend statement:

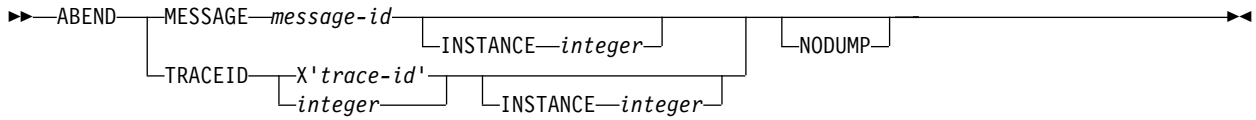

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

END Ends DIAGNOSE processing.

TYPE(*integer***, ...)**
Specifies one or more types of diagnose you wish to perform.

integer is the number of types of diagnoses. The maximum number of types is 32. An IBM Service Representative defines the types as needed to diagnose problems with IBM utilities.

ALLDUMPS(X'*abend-code***', ...)**
Forces a dump to be taken from any utility abend code.

X'*abend-code*' is a member of a list of abend codes to which the scope of ALLDUMPS is limited.

abend-code is a hexadecimal value.

NODUMPS(X'*abend-code***', ...)**
Suppresses the dump for any utility abend code.

X'*abend-code*' is a member of a list of abend codes to which the scope of NODUMPS is limited.

abend-code is a hexadecimal value.

DISPLAY
Formats the specified database items using SYSPRINT.

OBD *database-name.table-space-name*
Formats the object descriptor (OBD) of the table space.

database-name is the name of the database in which the table space belongs.

table-space-name is the name of the table space whose OBD is to be formatted.

ALL Formats all OBDs of the table space. The OBD of any relationship associated with the table space is also formatted.

TABLES
Formats the OBDs of all tables in the specified table spaces.

INDEXES
Formats the OBDs of all indexes in the specified table spaces.

SYSUTIL
Formats every SYSUTIL record.

MEPL
Dumps the module entry point lists (MEPLs) to SYSPRINT.

AVAILABLE
Displays the utilities installed on this subsystem in both bitmap and readable format. The presence or absence of utility products 5655-E62, 5655-E63, and 5697-E98 will affect the results of this display. See message DSNU862I for the output of this display.

DBET
Dumps the contents of a database exception table (DBET) to SYSPRINT. This option is intended to be used only under the direction of your IBM Support Center.

DATABASE *database-name*
Dumps the DBET entry associated with the specified database.

database-name is the name of the database.

TABLESPACE *database-name.table-space-name*
Dumps the DBET entry associated with the specified table space.

database-name is the name of the database.

table-space-name is the name of the table space.

INDEX *creator-name.index-name*
Dumps the DBET entry associated with the specified index.

creator-name is the ID of the creator of the index.

index-name is the name of the index.

WAIT After encountering the specified utility message or utility trace ID, a message is issued to the console and utility execution is suspended until the operator replies to that message, the utility job times out, or the utility job is canceled. This allows events to be synchronized while diagnosing concurrency problems. The utility waits for the operator to reply to the message, allowing the opportunity to time or synchronize events.

If neither the utility message nor the trace ID are encountered, WAIT processing continues.

ABEND
After encountering the specified utility message or utility trace ID, an abend is forced during utility execution.

If neither the utility message nor the trace ID are encountered, ABEND processing continues.

NODUMP
Suppresses the dump generated by an abend of DIAGNOSE.

MESSAGE *message-id*
Specifies a DSNU*xxx* or DSNU*xxxx* message that causes a wait or an abend to occur when that message is issued. Valid message IDs can be found in Part 2 of *DB2 Messages and Codes*.

message-id is the message in the form of U*xxx* or U*xxxx*.

INSTANCE *integer*
Specifies that a wait or an abend occurs when the MESSAGE option message has been encountered a specified number of times. If INSTANCE is not specified, a wait or abend occurs each time the message is encountered.

integer is the number of times a message is to be encountered before a wait or an abend occurs.

TRACEID *trace-id*
Specifies a trace ID that causes a wait or an abend to occur when the ID is encountered. Valid trace IDs can be found in data set *prefix*.SDSNSAMP(DSNWEIDS).

trace-id is a trace ID associated with the utility trace (RMID21), and can be specified in either decimal (*integer*) or hexadecimal (X'*trace-id*').

INSTANCE *integer*
Specifies that a wait or an abend occurs when the TRACEID option has been encountered a specified number of times. If INSTANCE is not specified, a wait or abend occurs each time the trace ID is encountered.

integer is the number of times a trace ID is to be encountered before a wait or an abend occurs.

Instructions for running DIAGNOSE

To run DIAGNOSE, you must:

1. Prepare the necessary data sets, as described in "Data sets used by DIAGNOSE".
2. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for DIAGNOSE, see "Sample control statements" on page 128.)
3. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 128. (For a complete description of the syntax and options for DIAGNOSE, see "Syntax and options of the control statement" on page 123.)
4. Run DIAGNOSE.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Data sets used by DIAGNOSE

Table 18 describes the data sets used by DIAGNOSE. Include statements in your JCL for each required data set.

Table 18. Data sets used by DIAGNOSE

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes

The following objects are named in the utility control statement and do not require DD cards in the JCL:

Database
Database to gather diagnosis information about.

Table space
Table space to gather diagnosis information about.

Index space
Index to gather diagnosis information about.

Instructions for specific tasks

To perform the following task, specify the options and values for those tasks in your utility control statement.

Forcing a utility abend with DIAGNOSE

DIAGNOSE can force an utility to abend when a specific message is issued. To force an abend when unique index or referential constraint violations are detected, you must specify the message that is issued when the error is encountered by using the MESSAGE option of the ABEND statement.

Instead of using a message, you can use the TRACEID option of the ABEND statement to specify a trace IFCID associated with the utility to force an abend.

Use the INSTANCE keyword to specify the number of times the specified message or trace record is generated before the utility abends.

Terminating or restarting DIAGNOSE

You can terminate DIAGNOSE with the TERM UTILITY command.

Concurrency and compatibility

DIAGNOSE can run concurrently on the same target object with any SQL operation or utility, except a utility running on DSNDB01.SYSUTILX.

Sample control statements

Example 1: Sample JCL for DIAGNOSE.

```
//STEP1    EXEC DSNUPROC,UID='IUJMU116.COPY1',
//         UTPROC='',
//         SYSTEM='V71A'
//SYSCOPY1  DD DSN=IUJMU116.COPY.STEP1.SYSCOPY1,DISP=(NEW,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
DIAGNOSE  ABEND MESSAGE U400
                 INSTANCE 1
                 NODUMP
                    COPY TABLESPACE DSN8D71A.DSN8S71E
                      COPYDDN SYSCOPY1
DIAGNOSE
          END
//*
```

Example 2: Force dump for utility abend. Force a dump for any utility abend that occurs during the execution of the COPY utility.

```
DIAGNOSE
 ALLDUMPS
 COPY TABLESPACE DSNDB06.SYSDBASE
DIAGNOSE END
```

Example 3: Force utility abend if message is issued. Abend the LOAD utility the fifth time message DSNU311 is issued. Do not generate a dump.

```
DIAGNOSE
 ABEND MESSAGE U311 INSTANCE 5 NODUMP
LOAD DATA RESUME NO
  INTO TABLE TABLE1
   (NAME POSITION(1) CHAR(20))
DIAGNOSE END
```

Example 4: Display SYSUTIL table. Display all rows in the SYSUTIL table, and the DB2 and utility MEPLs.

```
DIAGNOSE
 DISPLAY SYSUTIL
DIAGNOSE
 DISPLAY MEPL
```

Example 5: Abend LOAD utility if message is issued. Abend the LOAD utility when unique index key violations occur.

```
DIAGNOSE
 ABEND MESSAGE U344
```

Example 6: Force dump if abend with specified reason code. Cause a dump to be taken if an abend occurs with either of the specified reason codes.

```
DIAGNOSE
 ALLDUMPS(X'00E40322',X'00E40323')
```

Example 7: Display installed utilities. Display a bitmap and list of all installed DB2 utilities.

```
DIAGNOSE DISPLAY AVAILABLE
```

The output from this utility should look similar to this:

```
DSNU050I     DSNUGUTC - DIAGNOSE DISPLAY AVAILABLE
DSNU862I     DSNUDIAG - DISPLAY AVAILABLE UTILITIES.
                  MAP:  11111111111111111110000000000000
-------------------------------------------------------------------------------
|CATMAINT |CHECK     |COPY      |DIAGNOSE |LISTDEF  |LOAD      |MERGECOPY|MODIFY   |
|OPTIONS  |QUIESCE   |REBUILD   |RECOVER  |REORG    |REPAIR    |REPORT   |RUNSTATS |
|STOSPACE |TEMPLATE  |UNLOAD    |         |         |          |         |         |
|         |          |          |         |         |          |         |         |
-------------------------------------------------------------------------------
```

The MAP shows a ″1″ for each installed utility. Each position represents a specific utility. This example shows that the core utilities and DB2 Utilities Suite are installed.

Chapter 12. EXEC SQL

The EXEC SQL utility control statement declares cursors or executes dynamic SQL statements as part of the DB2 family cross loader function.

Output: The EXEC SQL control statement produces a result table when you specify a cursor.

Authorization required: The EXEC SQL statement itself requires no privileges to execute. The authorization rules defined for the dynamic preparation of the SQL statement specified by EXECUTE IMMEDIATE apply. See *DB2 SQL Reference* for authorization rules for each SQL statement.

Execution phases of EXEC SQL: The EXEC SQL control statement executes entirely in the EXEC phase. The EXEC phase can be restarted if necessary.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

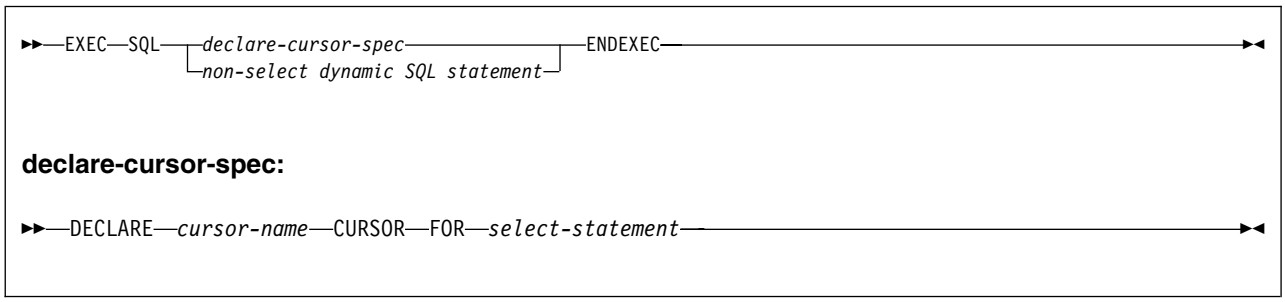

Option descriptions

For a description of parsing utility statements, and interpreting option identifiers and specifications, see "Control statement coding rules" on page 23.

non-select dynamic SQL statement
Specify a dynamic SQL statement that is used as input to EXECUTE IMMEDIATE. See Table 19 for a list of dynamic SQL statements that you can specify in a utility statement:

Table 19. Dynamic SQL statements you can specify in a utility statement

ALTER	RENAME
COMMENT ON	REVOKE
COMMIT	ROLLBACK
CREATE	SET CURRENT DEGREE

© Copyright IBM Corp. 1983, 2001

Table 19. Dynamic SQL statements you can specify in a utility statement (continued)

DELETE	SET CURRENT LOCALE LC_CTYPE
DROP	SET CURRENT OPTIMIZATION HINT
EXPLAIN	SET PATH
GRANT	SET CURRENT PRECISION
INSERT	SET CURRENT RULES
LABEL ON	SET CURRENT SQLID
LOCK TABLE	UPDATE

Each SQL statement runs as a separate thread. When the SQL statement is executed, the specified statement string is parsed and checked for errors. If the SQL statement is invalid, it is not executed and the error condition is reported. If the SQL statement is valid, but an error occurs during execution that error condition is reported. When an error occurs, the utility terminates.

cursor-name Specifies the cursor name. The name must not identify a cursor that is already declared. If the cursor is used to load data from a remote location, the table that is specified in the SQL-statement must have a three-part name.

select-statement
Specify the result table for the cursor. See *DB2 SQL Reference* for a description of the select-statement.

Concurrency and compatibility

The EXEC SQL control statement can be used with any utility that allows concurrent SQL access on a table space. Other databases are not effected.

Sample control statements

Example 1: Create table: Create table MYEMP like sample table EMP.

```
EXEC SQL
  CREATE TABLE MYEMP LIKE DSN8710.EMP
  CCSID EBCDIC
ENDEXEC
```

Example 2: Insert into table: Insert all rows from sample table EMP into table MYEMP.

```
EXEC SQL
  INSERT INTO MYEMP SELECT * FROM DSN8710.EMP
ENDEXEC
```

Example 3: Declare a cursor: Declare a cursor that selects all the rows and columns from sample table EMP.

```
EXEC SQL
  DECLARE C1 CURSOR FOR SELECT * FROM DSN8710.EMP
ENDEXEC
```

Example 4: Use a mapping table: Create a mapping table and mapping table index. Run online REORG and drop the mapping table.

```
EXEC SQL
  CREATE TABLE MYMAPPING_TABLE
    (TYPE         CHAR( 01 ) NOT NULL,
     SOURCE_RID   CHAR( 05 ) NOT NULL,
     TARGET_XRID  CHAR( 09 ) NOT NULL,
     LRSN         CHAR( 06 ) NOT NULL)
   IN DSN8D71P.DSN8S71Q
   CCSID EBCDIC
ENDEXEC

EXEC SQL
  CREATE UNIQUE INDEX MYMAPPING_INDEX
   ON MYMAPPING_TABLE
    (SOURCE_RID ASC,
     TYPE,
     TARGET_XRID,
     LRSN)
   USING STOGROUP DSN8G710
   PRIQTY 120 SECQTY 20
   ERASE NO
   BUFFERPOOL BP0
   CLOSE NO
ENDEXEC

REORG TABLESPACE DSN8D71P.DSN8S71C
 COPYDDN(COPYDDN)
 SHRLEVEL CHANGE
 DEADLINE CURRENT_TIMESTAMP+8 HOURS
 MAPPINGTABLE MYMAPPING_TABLE
 MAXRO 240 LONGLOG DRAIN DELAY 900
 SORTDEVT SYSDA SORTNUM 4
 STATISTICS TABLESPACE DSN8D71P.DSN8S71C
            INDEX(ALL)

EXEC SQL
  DROP TABLE MYMAPPING_TABLE
ENDEXEC
```

Chapter 13. LISTDEF

Traditional database administration includes querying the catalog to identify objects to be processed. After they are identified, utility jobs must be altered to specify the correct utility control statements, DD names and JCL DD cards. DB2 for OS/390 and z/OS Version 7 lets you group database objects into reusable lists and to dynamically allocate the utility data sets that are required to process that list. These new capabilities are the most useful in installations in which database objects have rigid, meaningful naming conventions that can be used to group objects for administration purposes, and where rigid utility data set naming conventions have been adopted.

DB2 utilities are generally invoked against one or more database objects: either a table space or an index space. Operating against multiple objects generally requires a user to run multiple utility invocations or to specify an itemized list of objects on a single invocation. DB2 now applies a pattern-matching expression to generate a list of objects, which it can then pass to one or more specific utilities for processing.

The LISTDEF control statement defines a list of objects and assigns a name to the list. A complete LISTDEF control statement must include:

- The name of the list
- An INCLUDE clause, optionally followed by additional INCLUDE or EXCLUDE clauses to either include or exclude objects from the list
- The type of objects the list contains, either TABLESPACES or INDEXSPACES
- The object to be used in the initial catalog lookup

  The search for objects can begin with databases, table spaces, index spaces, tables, indexes or other lists. The resulting list will only contain TABLESPACES, INDEXSPACES, or both.

In the LISTDEF control statement, you can include a list of specific objects, a generic pattern-matching expression, or both. Generic names contain pattern-matching characters that are similar to the LIKE predicate in the SQL language (See Table 23 on page 145 for a list of characters). Generic names are evaluated at execution time from information stored in the DB2 catalog. All-inclusive lists, such as `DATABASE *` or `TABLESPACE *.*` are not supported.

If you specify DB2 directory objects (DSNDB01) and DB2 catalog objects (DSNDB06) in object lists, you must use the fully qualified TABLESPACE or INDEXSPACE names for those objects. Pattern-matching is not supported for catalog or directory objects. DB2 issues error messages for any catalog or directory objects that are invalid for a utility. For more information, see “Specifying catalog and directory objects in LISTDEFs” on page 147.

Optionally, you can add related objects to the list by specifying keywords that indicate a relationship, such as referentially related objects or auxiliary related objects. Valid specifications include:

- RI (related by referential constraint)
- BASE (non-LOB objects)
- LOB (LOB objects)
- ALL (both BASE and LOB objects)
- TABLESPACES (related table spaces)
- INDEXSPACES (related index spaces)

© Copyright IBM Corp. 1983, 2001

Keywords perform two functions: They determine which objects are related, and they filter the contents of the list. The behavior of these keywords varies depending on the type of object you start with. For example, if your initial object is a LOB object, then the LOB keyword is ignored. If, however, the initial object is not a LOB object, then the LOB keyword determines which LOB objects are related, and filters non-LOB objects out of the list.

You can use the generated list for subsequent utility invocations.

Restriction: Objects created with the DEFINE NO attribute are excluded from all LISTDEF lists.

Output: Output from the LISTDEF control statement consists of a list with a name.

Authorization required:

Sample control statements

To execute the LISTDEF control statement, you must have SELECT authority on SYSIBM.SYSINDEXES, SYSIBM.SYSTABLES and SYSIBM.SYSTABLESPACE.

Additionally, you must have the authority to execute the utility which is used to process the list, as currently documented in the “Authorization required” section of each utility in this book.

Execution phases of LISTDEF: The LISTDEF control statement executes entirely within the UTILINIT phase.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

In addition, you can create a library of LISTDEF control statements. The default *ddname* of this library data set is SYSLISTD; you can change the DD name by using the OPTIONS control statement. The library contains only LISTDEF control statements. If a LISTDEF of the same name is found in SYSIN DD statement and in SYSLISTD data sets, DB2 uses the name from SYSIN.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

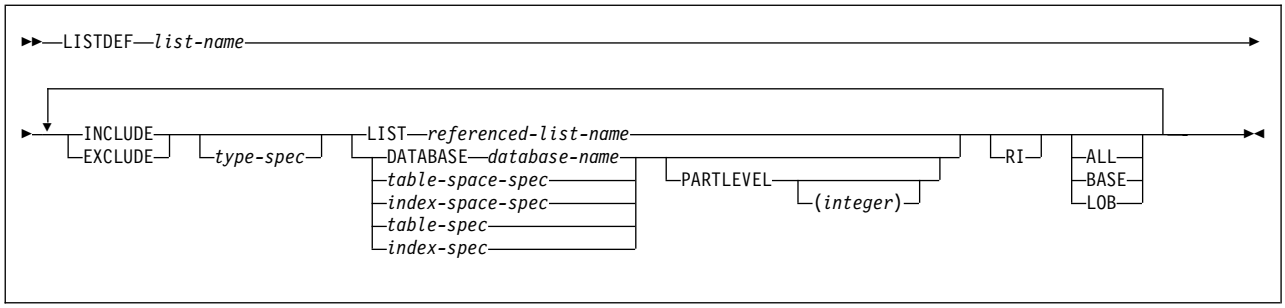

type-spec:

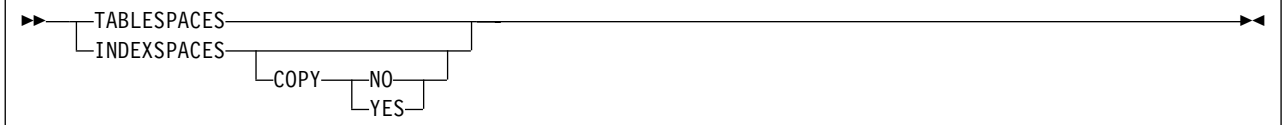

table-space-spec:

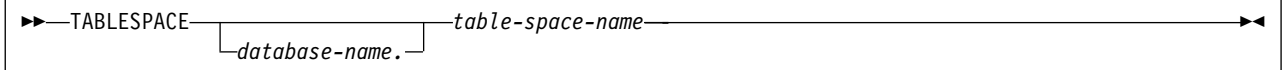

index-space-spec:

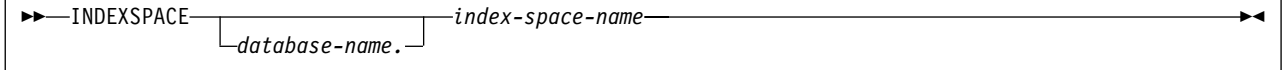

table-spec:

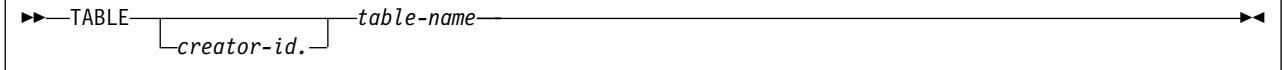

index-spec:

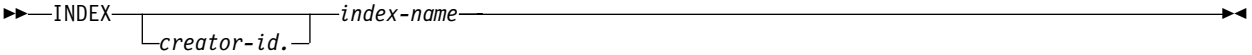

Option descriptions

For a description of how to parse utility statements, and how to read a list of option identifiers and specifications, see "Control statement coding rules" on page 23.

LISTDEF *list-name*

Defines a list of DB2 objects and assigns a name to the list. The list name makes the list available for subsequent execution as the object of a utility control statement or as an element of another LISTDEF.

list-name is the name (up to 18 alphanumeric characters in length) assigned to the defined list.

You can put LISTDEF statements either in a separate LISTDEF library data set or before a DB2 utility control statement that refers to the *list-name*. The default DD name for a LISTDEF library data set is SYSLISTD. You can alter the DD name by using the OPTIONS LISTDEFDD control statement.

INCLUDE Specifies that the list of objects that results from the expression that follows is to be added to the list. You can first specify an INCLUDE clause. You can specify subsequent INCLUDE or EXCLUDE clauses in any order to add to or delete clauses from the existing list. In general, INCLUDE and EXCLUDE clauses are processed in the order specified. However, some utilities alter the list order for optimal processing:

- CHECK INDEX, REBUILD and RUNSTATS INDEX process all index spaces that are related to a given table space at one time, regardless of list order.
- UNLOAD processes all specified partitions of a given table space at one time regardless of list order.

You must specify either INCLUDE or EXCLUDE. There is no default specification.

DB2 constructs the list, one specification at a time, by adding objects to or removing objects from the list. If an EXCLUDE attempts to remove an object that is not yet on the list, DB2 ignores the EXCLUDE clause of that object and proceeds to the next INCLUDE or EXCLUDE clause. Be aware that a subsequent INCLUDE can return the excluded object to the list.

For detailed information about the order of INCLUDE and EXCLUDE processing, see "Key LISTDEF operations" on page 146.

EXCLUDE After the initial INCLUDE clause, EXCLUDE specifies that the list of objects that results from the expression that follows is to be excluded from the list if the objects are on the list. If the objects are

not on the list, they are ignored, and DB2 proceeds to the next INCLUDE or EXCLUDE clause. You can specify subsequent INCLUDE or EXCLUDE clauses to add objects to or exclude additional objects from the list.

You must specify either INCLUDE or EXCLUDE. There is no default specification.

DB2 processes LISTDEF INCLUDE and EXCLUDE clauses in the order in which they are specified. DB2 constructs a list, one specification at a time, by adding objects to or removing objects from the list. If an EXCLUDE attempts to remove an object clause that is not yet on the list, DB2 ignores the EXCLUDE of that object and proceeds to the next INCLUDE or EXCLUDE clause. Be aware that a subsequent INCLUDE clause can return the excluded object to the list.

For detailed information about the order of INCLUDE and EXCLUDE processing, see "Key LISTDEF operations" on page 146.

TABLESPACES

Specifies that the INCLUDE or EXCLUDE object expression creates a list of related TABLESPACES.

TABLESPACES is the default for lists developed from *table-space-spec* and *table-spec* options.

There is no default value for lists developed from a *list-spec* option. The LISTDEF referred to by the *list-spec* is used unless you specify TABLESPACES or INDEXSPACES. There is no default value for lists created from the database-spec option.

The result of the TABLESPACES keyword varies depending on the type of object you specify in the INCLUDE or EXCLUDE clause.

Table 20. TABLESPACES keyword by object type result table

DATABASE	Returns all table spaces contained within the database
TABLESPACE	Ignored
TABLE	Returns the table space containing the table
INDEXSPACE	Returns the table space containing the related table
INDEX	Returns the table space containing the related table
LIST	Returns the table spaces from the expanded referenced list

INDEXSPACES

Specifies that the INCLUDE or EXCLUDE object expression creates a list of related INDEXSPACES.

INDEXSPACES is the default for lists developed from *index-space-spec* and *index-spec*.

There is no default value for lists developed from a *list-spec*. The LISTDEF referred to by the *list-spec* is used unless you specify TABLESPACES or INDEXSPACES. There is no default value for lists developed from database-spec.

The result of the INDEXSPACES keyword varies depending on the type of object you specify in the INCLUDE or EXCLUDE clause.

Table 21. INDEXSPACES keyword by object type result table

DATABASE	Returns all index spaces contained within the database
TABLESPACE	Returns all index spaces for indexes over all tables in the table space
TABLE	Returns all index spaces for indexes over the table
INDEXSPACE	Ignored
INDEX	Returns the index space containing the index
LIST	Returns the index spaces from the expanded referenced list

DATABASE *database-name*
Specifies the database that is used for the initial search for the object.

You can specify the database-name explicitly or as a pattern-matched name. `DATABASE *` and `DATABASE %` are not supported.

If you specify DATABASE, you must also specify either TABLESPACES or INDEXSPACES. DB2 includes all table spaces or index spaces in *database-name* that satisfy the pattern-matching expression in the list.

You cannot specify DSNDB01, DSNDB06, DSNDB07 or user defined work file databases in a LISTDEF.

TABLESPACE *database-name.table-space-name*
Specifies the table space that is used for the initial search for the object.

If you specify TABLESPACE, the default list type is TABLESPACES. All table spaces that satisfy the pattern-matching expression are included in the list unless modified by other keywords. `TABLESPACE *.*` and `TABLESPACE %.%` are not supported.

database-name
Specifies the name of the database to which the table space belongs.

The **default** is **DSNDB04**.

table-space-name
Specifies the name of the table space.

You can explicitly specify or use pattern-matching characters to specify *database-name* or *table-space-name*, or both.

You can not include any objects in DSNDB07 or any user defined work file databases in a LISTDEF. Pattern matching is not supported for DSNDB01 and DSNDB06 objects.

INDEXSPACE *database-name.index-space-name*
Specifies the index space to be used for the initial search for the object.

If you specify INDEXSPACE, the default list type is INDEXSPACES. All index spaces that satisfy the pattern-matching expression are

included in the list unless the index spaces are excluded by other LISTDEF options. `INDEXSPACE *.*` and `INDEXSPACE %.%` are not supported.

database-name
Specifies the name of the database the index space belongs to.

The **default** is **DSNDB04**.

index-space-name
Specifies the name of the index space.

You can explicitly specify or use pattern-matching characters to specify *database-name* or *index-space-name*, or both.

You can not include any objects in DSNDB07 or any user defined work file databases in a LISTDEF. Pattern-matching is not supported for DSNDB01 and DSNDB06 objects.

TABLE *creator-id.table-name*
Specifies the table to be used for the initial search for the object.

If you specify TABLE, the default list type is TABLESPACES. All table spaces that contain tables that satisfy the pattern-matching expression are included in the list unless modified by other keywords. `TABLE *.*` and `TABLE %.%` are not supported.

creator-id Specifies the qualifier creator ID for the table.

The **default** is the user identifier for the utility.

table-name Specifies the name of the table.

You can explicitly specify or use pattern-matching characters to specify *creator-id* or *table-name*, or both. However, the underscore pattern-matching character is ignored in a table name.

Pattern-matching is not supported for catalog and directory objects. In a LISTDEF statement, you must include catalog and directory objects by their fully qualified names.

INDEX *creator-id.index-name*
Specifies the index to be used for the initial search for the object.

If you specify INDEX, the default list type is INDEXSPACES. All index spaces that contain indexes that satisfy the pattern-matching expression are included in the list unless modified by other keywords. `INDEX *.*` and `INDEX %.%` are not supported.

creator-id Specifies the qualifier creator ID for the index.

The **default** is the user identifier for the utility.

index-name Specifies the name of the index.

You can explicitly specify or use pattern-matching characters to specify *creator-id* or *index-name*, or both. However, the underscore pattern-matching character is ignored in an index name.

Pattern-matching is not supported for catalog and directory objects. In a LISTDEF statement, you must include catalog and directory objects by their fully qualified names.

LIST *referenced-list-name*
Specifies the name of a LISTDEF defined object list that is to be expanded and used for the initial search for the object.

You must explicitly specify the *referenced-list-name*. Pattern-matching characters (%,*, ?, and _) are not allowed.

There is no default type value for lists developed from a *list-spec* option. The list is expanded as defined, and then modified by subsequent keywords. You can specify a list type of TABLESPACES or INDEXSPACES to apply an additional constraint to the existing list.

The LIST keyword is useful for making aggregate lists of lists, for excluding entire lists from other lists, and for developing lists of objects that are related to other lists.

PARTLEVEL Specifies the partition granularity for partitioned table spaces and partitioning indexes that are contained in the list. PARTLEVEL is ignored for non-partitioned objects. You cannot specify the PARTLEVEL with the RI keyword.

(*n***)** *n* is the integer partition number where *n*>0. If you specify *n*, an entry is added to the list for the *n*th partition of each partitioned object.

If you specify PARTLEVEL without (*n*), the resulting list contains one entry for each partition in the partitioned object.

An INCLUDE with PARTLEVEL can only be removed from the list by an EXCLUDE with PARTLEVEL.

The following example removes partition 12 because the INCLUDE and the EXCLUDE both specify PARTLEVEL:

```
LISTDEF INCLUDE TABLESPACE X PARTLEVEL
        EXCLUDE TABLESPACE X PARTLEVEL(12)
```

The EXCLUDE clause in both of the following examples below does not affect the list contents because the PARTLEVEL specifications are mismatched:

```
LISTDEF INCLUDE TABLESPACE X
        EXCLUDE TABLESPACE X PARTLEVEL(12)

LISTDEF INCLUDE TABLESPACE X PARTLEVEL(12)
        EXCLUDE TABLESPACE X
```

COPY Specifies whether indexes that were defined with or altered to COPY YES or COPY NO attributes are to be included or excluded in this portion of the list. If you omit COPY, all index spaces that satisfy the INCLUDE or EXCLUDE expression, regardless of their COPY attribute, are included or excluded in this portion of the list. If specified, this keyword must immediately follow the INDEXSPACES keyword. If specified elsewhere, the keyword COPY is interpreted as the start of the COPY utility control statement.

YES Only index spaces that were defined with or altered to COPY YES will be included in this portion of the list. Use

INCLUDE with COPY YES to develop a list of index spaces that the COPY utility can process.

NO Only index spaces that were defined with or altered to COPY NO will be included in this portion of the list. Use EXCLUDE with COPY NO to remove indexes the COPY utility can not process from a larger list.

RI Specifies that all objects that are referentially related to the object expression (PRIMARY KEY <--> FOREIGN KEY) are to be included in the list. Referential relationships are followed repeatedly until the entire referential set is developed.You cannot specify RI with PARTLEVEL(*n*).

LOB indicator keywords: Use one of three LOB indicator keywords to direct LISTDEF processing to follow auxiliary relationships to include related LOB objects in the list. The auxiliary relationship may be followed in either direction. LOB objects include the LOB table spaces, auxiliary tables, indexes on auxiliary tables and their containing index spaces.

There is no default LOB indicator keyword. If you do not specify BASE, LOB or ALL, DB2 will not follow the auxiliary relationships and will not filter LOB from BASE objects in the enumerated list.

ALL Specifies that both related BASE and LOB objects are to be included in the list. Auxiliary relationships will be followed from all objects resulting from the initial object lookup and both BASE and LOB objects will remain in the final enumerated list.

BASE Specifies that only base table spaces (non-LOB) and index spaces are to be included in this element of the list.

If the initial search for the object results in a base object, auxiliary relationships are not followed. If the initial search for the object results in a LOB object, the auxiliary relationship is applied to the base table space or index space, and only those objects become part of the resulting list.

LOB Specifies that only LOB table spaces and related index spaces containing indexes on auxiliary tables are to be included in this element of the list.

If the initial search for the object results in a LOB object, auxiliary relationships are not followed. If the initial search for the object results in a base object, the auxiliary relationship is applied to the LOB table space or index space, and only those objects become part of the resulting list.

Instructions for using LISTDEF

To use a list that has been defined with the LISTDEF control statement for a specific utility, refer to the list name, prefixed by the LIST keyword. For example, you could QUIESCE all objects in a list by specifying:

```
QUIESCE LIST X
```

The LIST keyword is supported by the following utilities. When possible, utility processing optimizes the order of list processing as indicated:

Table 22. Pattern-matching character comparison table

CHECK INDEX	items grouped by related table space
COPY	items processed in order specified on a single call to COPY and PARALLEL keyword is supported
COPYTOCOPY	items processed in order specified on a single call to COPYTOCOPY
MERGECOPY	items processed in order specified
MODIFY RECOVERY	items processed in order specified
MODIFY STATISTICS	items processed in order specified
QUIESCE	all items processed in order specified on a single call to QUIESCE
REBUILD	items grouped by related table space
RECOVER	items processed in order specified on a single call to RECOVER
REORG	items processed in order specified
REPORT	items processed in order specified
RUNSTATS INDEX	items grouped by related table space
RUNSTATS TABLESPACE	items processed in order specified

Some utilities such as COPY and RECOVER, can process a LIST without a specified object type. Object types are determined from the list contents. Other utilities, such as QUIESCE and REORG INDEX, must know the object type to be processed before processing can begin. These utilities require that you specify an object type in addition to the LIST keyword, for example: QUIESCE LIST and REORG INDEX LIST. See the syntax diagrams for an individual utility for details.

Many utilities require output data sets to process an object. In those cases, you should use the TEMPLATE control statement to specify the naming convention and, optionally, the allocation parameters for each output data set type. Templates, like lists, can be reused as long as the naming convention is robust enough to prevent duplicate data set names from being allocated. (See “Chapter 28. TEMPLATE” on page 463 for more details). In some cases it is possible to use traditional JCL DD cards with LISTDEF lists but this is usually not practical. Only when small lists are being processed one object at a time is this practical.

Together, the LISTDEF and TEMPLATE facilities allow faster development of utility job streams, and require fewer modifications when the underlying list of database objects change.

The OPTIONS PREVIEW control statement can be used to preview the list contents before actual processing. The OPTIONS ITEMERROR control statement can be used to alter the handling of errors which may occur during list processing. The OPTIONS LISTDEFDD control statement may be used to switch LISTDEF library data sets between control statements within a job step. The default is LISTDEFDD SYSLISTD.

Generating object lists using LISTDEF

The statement of a typical utility control statement consists of a utility name that is followed by one or more object specifications (object type and object name) that are followed by related parameters and options. The pattern-matching facility does not

change the basic control statement structure. Rather, it provides a mechanism for defining a *list* and assigning that list a name for use in subsequent utility invocations.

You can define a *list* by using a LISTDEF control statement. A LISTDEF control statement contains a list of specific objects, a generic pattern-matching expression, or both. You can use four special pattern-matching characters (%, *, _, ?) to define *generic* object names in a LISTDEF statement. Utility processes access the DB2 catalog at execution time and dynamically expand each generic object name into an equivalent *enumerated* list. A utility processes this enumerated list either sequentially or in parallel, depending on the utility function and the parameters that you specify.

Utilities do not support SYSUTILX related objects being within a LISTDEF specification. Pattern-matching of DB2 catalog and directory objects (DSNDB06 and DSNDB01) are also not supported. Catalog and directory objects must be included in a LISTDEF explicitly by their full TABLESPACE or INDEXSPACE name. Even if Catalog and Directory objects match a LISTDEF pattern they will not be included in the list. You cannot specify the following objects in a LISTDEF:

- TABLESPACE DSNDB01.SYSUTILX
- TABLE SYSIBM.SYSUTILX
- TABLE SYSIBM.SYSUTIL
- INDEXSPACE DSNDB01.DSNLUX01
- INDEXSPACE DSNDB01.DSNLUX02
- INDEX SYSIBM.DSNLUX01
- INDEX SYSIBM.DSNLUX02

To process those objects, you must use syntax from releases prior to Version 7.

Valid pattern-matching characters

Specify pattern-matching object names by using the following pattern-matching characters:

Table 23. LISTDEF pattern-matching characters

LISTDEF pattern-matching character	Equivalent symbol used in SQL LIKE predicates	Usage notes
Percent sign (%)	Percent sign (%)	Performs the same function.
Question mark (?)	Underscore (_)	Use the question mark (?) instead of underscore (_) as a pattern-matching character in table and index names. The underscore character (_) in table and index names represents a single occurrence of itself.
Asterisk (*)	Percent sign (%)	Performs the same function.
Underscore (_)	Underscore (_)	Use the underscore (_) as an alternative to the question mark (?) for database, table space, and index space names.

Important: The LISTDEF control statement does not support pattern-matching for either catalog or directory objects.

The following example illustrates how to specify a list of objects by using LISTDEF and pattern-matching characters:

```
LISTDEF PAYROLL INCLUDE TABLESPACE PAYROLL.*
                INCLUDE INDEXSPACE PAYROLL.*IX
                EXCLUDE TABLESPACE PAYROLL.TEMP*
                EXCLUDE INDEXSPACE PAYROLL.TMPIX*
COPY LIST PAYROLL ...
```

How to use LISTDEF specifications

Specify LISTDEF control statements in the SYSIN DD statement prior to the utility control statement that references it, or in one or more LISTDEF data sets. Specify the LISTDEF data set DD name on the OPTIONS control statement as LISTDEFDD ddname. This LISTDEF statement applies to all subsequent utility control statements until either the end of input or until you specify another OPTIONS LISTDEFDD ddname. The default LISTDEF DD name is SYSLISTD. A LISTDEF data set may contain only LISTDEF control statements. Any LISTDEF defined within the SYSIN DD statement overrides another LISTDEF definition of the same name found in a LISTDEF data set.

Database administrators that use LISTDEF can standardize object lists and the utility control statements that refer to them, which reduces the need to customize or alter utility job streams.

Key LISTDEF operations

You must first specify an INCLUDE clause in a LISTDEF statement. You can subsequently specify one or more INCLUDE or EXCLUDE clauses to refine the object list. Utility processes expands the initial INCLUDE list, and then sequentially expand each additional INCLUDE or EXCLUDE clause to add or remove specified objects from the list. Each successive INCLUDE or EXCLUDE within one LISTDEF is processed in top down order and adds objects to or removes objects from the most recent list.

For example, to generate a list of all table spaces in the ACCOUNT database but exclude all LOB table spaces, you might specify:

```
LISTDEF ACCNT INCLUDE TABLESPACES DATABASE ACCOUNT BASE
```

If you want a list of only LOB index spaces in a given database, you might specify:

```
LISTDEF ACLOBIX INCLUDE INDEXSPACES DATABASE ACCOUNT LOB
```

DB2 processes the keywords that constrain each element in the following order:

1. Perform the initial search for the object that is based on the specified pattern-matching expression, including PARTLEVEL specification if specified.
2. Add or remove related objects and filter the list elements based on the list type specified, either TABLESPACES or INDEXSPACES (COPY YES or COPY NO).
3. Add or remove related objects depending on the presence or absence of the RI, BASE, LOB, or ALL keywords.
4. The utility has either INCLUDED the resulting list in the existing list or EXCLUDED it from the list.
5. Process the next clause in the list.

Database administrators can check utility control statements without executing those statements by using the PREVIEW function. When you run a utility using the PREVIEW function, it expands a LISTDEF into the equivalent enumerated list, prints it to SYSPRINT, and stops execution. Specify PREVIEW in one of two ways, either as a JCL parameter or on the OPTIONS PREVIEW control statement.

For details about the OPTIONS PREVIEW function, see “Syntax and options of the control statement” on page 136.

Specifying catalog and directory objects in LISTDEFs

LISTDEF lets you explicitly specify DB2 catalog and directory table spaces and index spaces by their fully qualified names in object lists. Pattern-matching is not supported for catalog or directory objects.

While DB2 catalog and directory objects can appear in LISTDEF lists, these objects might be invalid for a utility and result in an error message. Valid examples of include/exclude clauses are:

- TABLESPACE DSNDB06.SYSDBASE
- TABLESPACES TABLESPACE DSNDB06.SYSDBASE
- INDEXSPACE DSNDB06.DSNDXX01
- INDEXSPACES INDEXSPACE DSNDB06.DSBDXX01

To eliminate catalog access during the processing of catalog and directory objects, other keywords that can initiate such access are not supported, specifically:

- DATABASE
- TABLE
- INDEX,
- BASE
- LOB
- ALL
- databases DSNDB01, DSNDB06 and DSNDB07
- table or indexes with creator-id of SYSIBM

All LISTDEF lists automatically exclude work file databases, consisting DSNDB07 objects and user-defined work file objects, because DB2 utilities do not process them.

Concurrency and compatibility

LISTDEF is a control statement that is used to set up an environment for another utility to follow. The LISTDEF list is stored until it is referenced by a specific utility. When referenced by an utility, the list expands. At that time, the concurrency and compatibility restrictions of that utility apply, with the additional restriction that the catalog tables necessary to expand the list must be available for read only access.

Sample control statements

Example 1: List all objects in the PAYROLL database.

```
LISTDEF EXAMPLE1 INCLUDE TABLESPACES DATABASE PAYROLL
                 INCLUDE INDEXSPACES DATABASE PAYROLL
```

Example 2: List table spaces that meet a name pattern, by partition.

```
LISTDEF EXAMPLE2 INCLUDE TABLESPACE PAY*.* PARTLEVEL
```

Example 3: List COPY YES indexes for the list in EXAMPLE2.

```
LISTDEF EXAMPLE3 INCLUDE LIST EXAMPLE2 INDEXSPACES COPY YES
```

Example 4: List all partitions of a table space, except for partition 10.

```
LISTDEF EXAMPLE4 INCLUDE TABLESPACE A.B PARTLEVEL
                 EXCLUDE TABLESPACE A.B PARTLEVEL(10)
```

Example 5: Build LISTDEF library data sets and then run the QUIESCE utility. In this example, the NAME1 list is stored in a sequential data set, and the NAME2 list is stored in a PDS data set. The LISTDSN dd card concatenates the previously defined data sets into a library data set that is used with the QUIESCE utility for the NAME3 list.

```
//CREATE1  JOB 'USER=NAME',CLASS=A,...
//*-------------------------------------------------------
//* Create an input data set.
//*-------------------------------------------------------
//LOAD1    EXEC PGM=IEBGENER
//SYSPRINT DD DUMMY
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=JULTU103.TCASE.DATA2,
//         DISP=(NEW,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND),
//         DCB=(RECFM=FB,LRECL=80,BLKSIZE=2400)
//SYSUT1   DD *
  LISTDEF NAME1 INCLUDE TABLESPACE DBLT0301.TLLT031A
                INCLUDE TABLESPACE DBLT0301.TSLT031B
/*


//CREATE2  JOB 'USER=NAME',CLASS=A,...
//*-------------------------------------------------------
//* Create an input data set.
//*-------------------------------------------------------
//CRECNTL  EXEC PGM=IEFBR14
//CNTL     DD  DSN=JULTU103.TCASE.DATA3,UNIT=SYSDA,
//             VOL=SER=SCR03,
//             SPACE=(TRK,(2,2,2)),DCB=(DSORG=PO,
//             LRECL=80,RECFM=FB,BLKSIZE=4560),
//             DISP=(NEW,CATLG,CATLG)
/*
//*-------------------------------------------------------
//* Create member of input data set.
//*-------------------------------------------------------
//FILLCNTL EXEC PGM=IEBUPDTE
//SYSPRINT DD  SYSOUT=*
//SYSUT1   DD  DSN=JULTU103.TCASE.DATA3,DISP=OLD,
//         UNIT=SYSDA,
//         VOL=SER=SCR03,DCB=(DSORG=PO,
//         LRECL=80,RECFM=FB,BLKSIZE=4560)
//SYSUT2   DD  DSN=JULTU103.TCASE.DATA3,DISP=OLD,
//         UNIT=SYSDA,
//         VOL=SER=SCR03,DCB=(DSORG=PO,
//         LRECL=80,RECFM=FB,BLKSIZE=4560)
//SYSIN  DD  DATA
./ ADD NAME=MEM1
  LISTDEF NAME2 INCLUDE TABLESPACE DBLT0302.TLLT032A
                INCLUDE TABLESPACE DBLT0302.TSLT032B
                INCLUDE TABLESPACE DBLT0302.TPLT032C
./ ENDUP
/*


//QUIESCE  JOB 'USER=NAME',CLASS=A,...
//*******************************************************
//* QUIESCE LISTDEF DD LILSTDEF datasets
//*******************************************************
//STEP1    EXEC DSNUPROC,UID='JULTU103.QUIESC2',
//            UTPROC='',SYSTEM='SSTR'
//LISTDSN  DD DSN=JULTU103.TCASE.DATA2,DISP=SHR
//         DD DSN=JULTU103.TCASE.DATA3(MEM1),DISP=SHR
```

```
//SYSUT1   DD DSN=JULTU103.QUIESC2.STEP1.SYSUT1,
//         DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD DSN=JULTU103.QUIESC2.STEP1.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
    OPTIONS LISTDEFDD LISTDSN
    LISTDEF NAME3 INCLUDE LIST NAME1
                  INCLUDE LIST NAME2
                  EXCLUDE TABLESPACE DBLT0302.TSLT032B
                  EXCLUDE TABLESPACE DBLT0301.TSLT031B
                  EXCLUDE TABLESPACE DBLT0302.TPLT032C
                  INCLUDE TABLESPACE DBLT0302.TSLT032B
    QUIESCE LIST NAME3
/*
```

(see “Sample control statements” on page 475 in “Chapter 28. TEMPLATE” on page 463 for additional examples showing LISTDEF and TEMPLATE used together)

Chapter 14. LOAD

Use LOAD to load one or more tables of a table space. LOAD loads records into the tables and builds or extends any indexes defined on them. If the table space already contains data, you can choose whether you want to add the new data to the existing data or replace the existing data. The loaded data is processed by any edit or validation routine associated with the table, and any field procedure associated with any column of the table.

For a diagram of LOAD syntax and a description of available options, see “Syntax and options of the control statement” on page 152. For detailed guidance on running this utility, see “Instructions for running LOAD” on page 182.

Output: Output from LOAD DATA consists of one or more of the following:

- A loaded table space or partition
- A discard file of rejected records
- A summary report of errors encountered during processing, generated only if you specify ENFORCE CONSTRAINTS or if the LOAD involves unique indexes

Related information: For information regarding ESA data compression, see Part 5 (Volume 2) of *DB2 Administration Guide*.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:

- Ownership of the table
- LOAD privilege for the database
- DBADM or DBCTRL authority for the database
- SYSCTRL or SYSADM authority

The following catalog tables can be loaded.

- SYSSTRINGS
- MODESELECT
- LUMODES
- LULIST
- USERNAMES
- LUNAMES
- LOCATIONS
- IPNAMES

You cannot run the LOAD utility on the DSNDB01 or DSNDB06 databases, except to add lines to SYSIBM.SYSSTRINGS.

LOAD operates on a table space, so you must have authority for all tables in the table space when you perform LOAD.

To run LOAD STATISTICS REPORT YES, the privilege set must include the SELECT privilege on the catalog tables.

Execution phases of LOAD: Table 24 on page 152 describes the phases of the LOAD utility.

© Copyright IBM Corp. 1983, 2001

Table 24. LOAD phases

Phase	Description
UTILINIT	Initialization and setup.
RELOAD	Loading of record types and writing of temporary file records for indexes and foreign keys. Check constraints are checked for each row. One pass through the sequential input data set is made. Internal commits are taken to provide commit points at which to restart in case operation should halt in this phase. Creates inline copies if you specified the COPYDDN or RECOVERYDDN keywords. If SORTKEYS is used, a subtask is started at the beginning of the RELOAD phase to handle the work of sorting the keys. The sort subtask initializes and waits for the main RELOAD phase to pass its keys to SORT. The RELOAD phase loads the data, extracts the keys, and passes them in memory for sorting. At the end of the RELOAD phase, the last key is passed to SORT, and record sorting completes. Note that load partition parallelism starts subtasks. PREFORMAT for table spaces occurs at the end of the RELOAD phase.
SORT	Sorting of temporary file records before creating indexes or validating referential constraints, if indexes or foreign keys exist. The SORT phase is skipped if all the following conditions apply for the data processed during the RELOAD phase: • There is not more than one key per table • All keys are the same type (index key, index foreign key, nonindexed foreign key) • The data being loaded or reloaded is in key order (if a key exists) • The data being loaded or reloaded is grouped by table and each input record is loaded into one table only. If you use the SORTKEYS keyword, SORT passes the sorted keys in memory to the BUILD phase, which builds the indexes.
BUILD	Creating indexes from temporary file records for all indexes defined on the loaded tables. Detection of duplicate keys. Preformatting of indexes occurs at the end of the build phase.
SORTBLD	If you specify a parallel index build, all activities that normally occur in both the SORT and BUILD phases occur in the SORTBLD phase instead.
INDEXVAL	Correction of unique index violations from the information in SYSERR, if any exist.
ENFORCE	Checking of referential constraints, and correction of violations. Information about violations of referential constraints are stored in SYSERR.
DISCARD	Copying of records causing errors from the input data set to the discard data set.
REPORT	Generation of a summary report, if you specified ENFORCE CONSTRAINT or if load index validation is performed. The report is sent to SYSPRINT.
UTILTERM	Cleanup.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

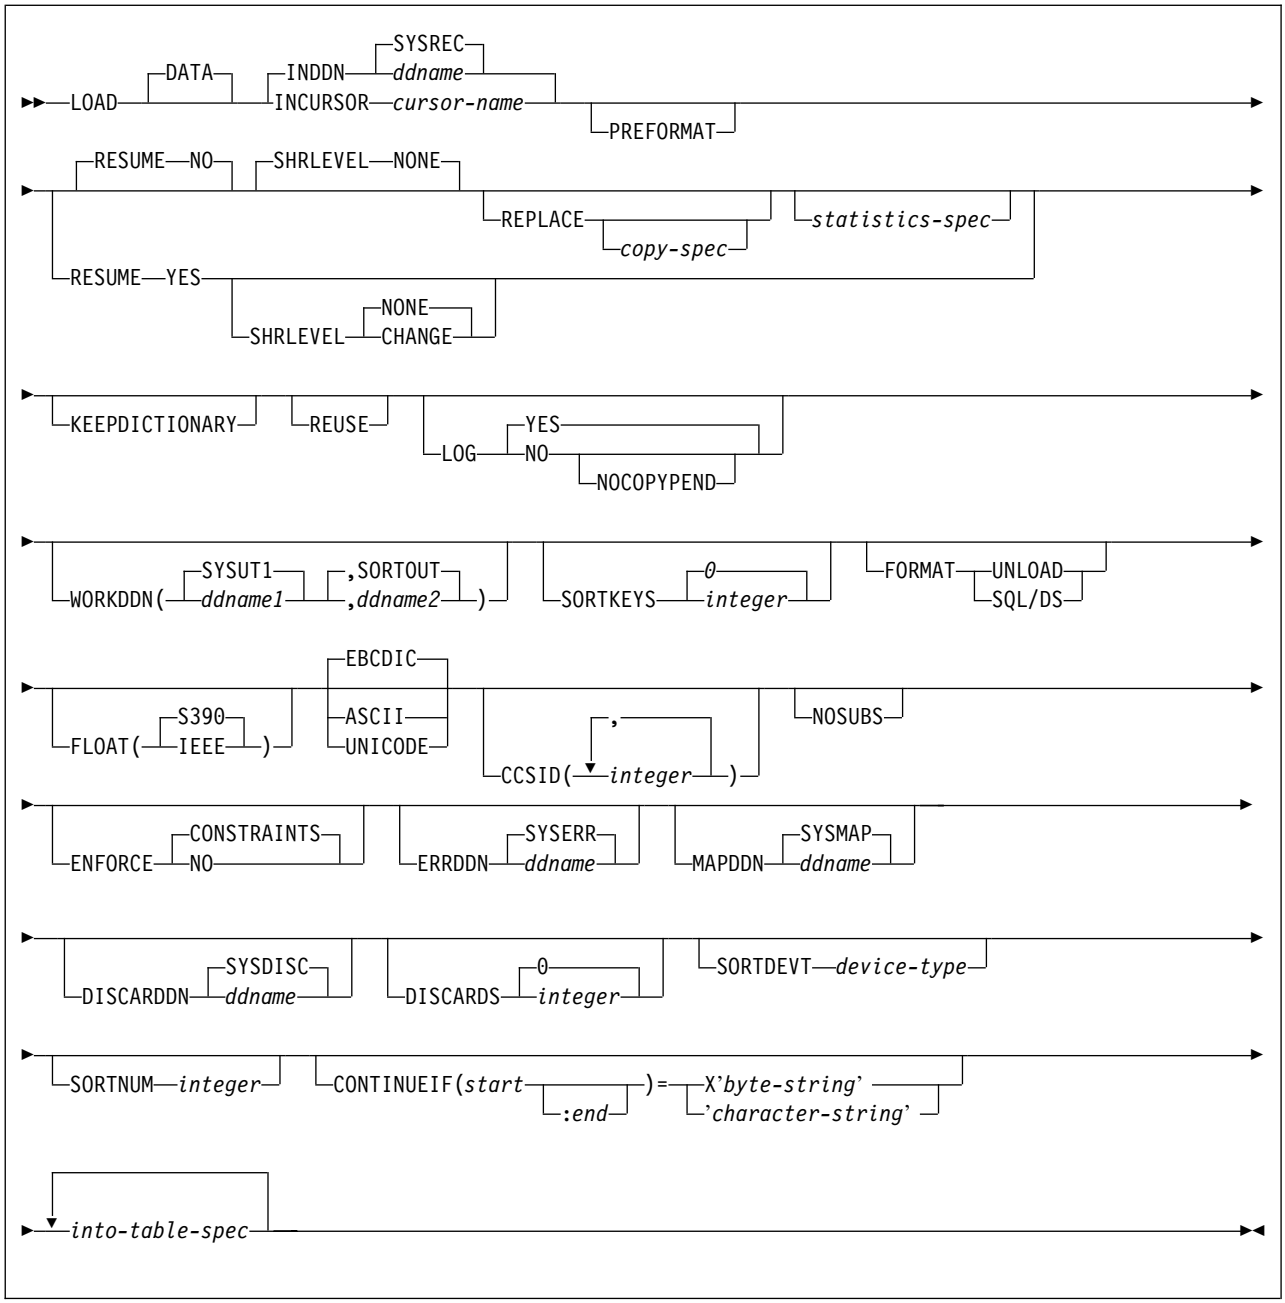

statistics-spec:

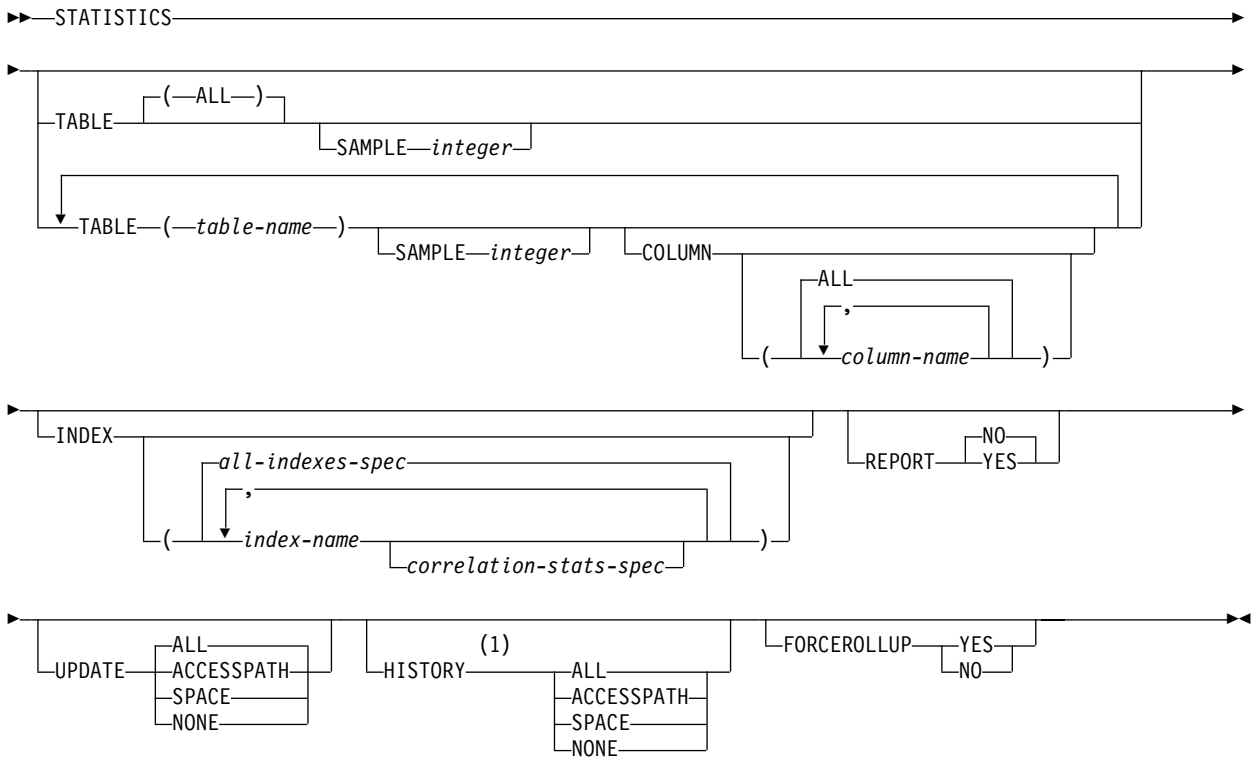

Notes:

1 What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

all-indexes-spec:

correlation-stats-spec:

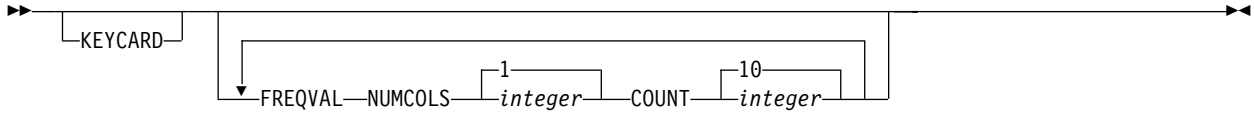

copy-spec:

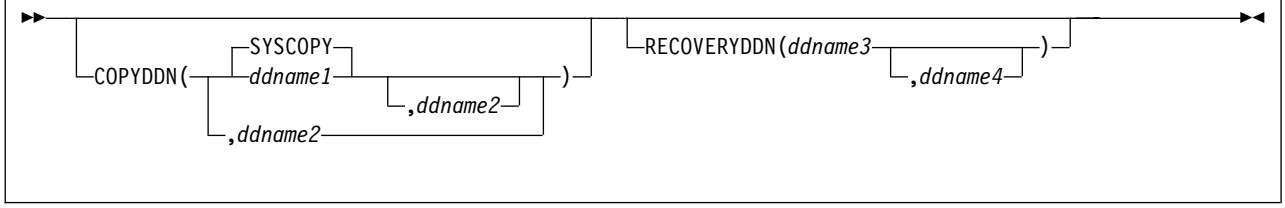

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

DATA Is used for clarity only. You identify the data selected for loading with *table-name* on the INTO TABLE option. See "Into-table-spec" on page 167 for a description of the statement.

INDDN *ddname*
Specifies the data definition (DD) statement or template identifying the input data set for the partition. The record format for the input data set must be fixed or variable. The data set must be readable by the basic sequential access method (BSAM).

The *ddname* is the name of the input data set.

The **default** is **SYSREC**.

INCURSOR *cursor-name*
Specifies the cursor for the input data set. The cursor must be declared before it is used by the LOAD utility. You cannot load into the same table where you defined the cursor. Use the EXEC SQL utility control statement to define the cursor for the DB2 family cross loader function. The INCURSOR option is incompatible with the SHRLEVEL CHANGE option.

cursor-name is the cursor name.

The column names in the SELECT statement must be identical to the column names in the table being loaded. The AS clause in the select list can be used to change the columns names returned by the SELECT statement.

PREFORMAT
Specifies that the remaining pages are preformatted up to the high allocated RBA in the table space and index spaces associated with the table specified in *table-name*. The preformatting occurs after the data has been loaded and the indexes are built.

PREFORMAT can operate on an entire table space and its index spaces, or on a partition of a partitioned table space and its corresponding partitioning index space. Specifying LOAD PREFORMAT (rather than PART *integer* PREFORMAT) tells LOAD to serialize at the table space level, which can inhibit concurrent processing of separate partitions. If you want to serialize at the partition level, specify PART *integer* PREFORMAT. See "Option descriptions for INTO TABLE" on page 169 for the description for specifying PREFORMAT at the partition level.

RESUME
Tells whether records are to be loaded into an empty or non-empty table space. For nonsegmented table spaces, space occupied by rows that have been marked as deleted or by rows of dropped tables is not reused.

Attention: Specifying LOAD RESUME (rather than PART *integer* RESUME) tells LOAD to drain the table space, which can inhibit concurrent processing of separate partitions. If you want to process other partitions concurrently, use "Into-table-spec" on page 167 to specify PART *integer* RESUME.

NO
Loads records into an empty table space. If the table space is not empty, and you have not used REPLACE, a message is issued and the utility job step terminates with a job step condition code of 8.

For nonsegmented table spaces containing deleted rows or rows of dropped tables, using the REPLACE keyword provides increased efficiency.

The **default** is **NO**, unless you override it with PART *integer* RESUME YES.

YES
Loads records into a non-empty table space. If the table space is empty, a warning message is issued, but the table space is loaded. Loading begins at the current end of data in the table space. Space occupied by rows marked as deleted or by rows of dropped tables is not reused.

SHRLEVEL
Specifies the extent to which applications can concurrently access the table space or partition during the LOAD utility. The parameter values below appear in order of increasing extent of allowed concurrent access.

NONE
Specifies that applications have no concurrent access to the table space or partition.

The **default** is **NONE**.

CHANGE
Specifies that applications can concurrently read and write the table space or partition it is loading. If you specify **SHRLEVEL CHANGE**, you can not specify the following parameters: **INCURSOR**, **RESUME NO**, **REPLACE**, **KEEPDICTIONARY**, **LOG NO**, **ENFORCE NO**, and **SORTKEYS**, **STATISTICS**, **COPYDDN**, **RECOVERYDDN**, **PREFORMAT**, **REUSE**, and **PART** integer **REPLACE**.

For a partition-directed **LOAD**, if you specify **SHRLEVEL CHANGE**, then only **RESUME YES** can be specified or inherited from the LOAD statement.

LOAD **SHRLEVEL CHANGE** does not perform the SORT, BUILD, SORTBLD, INDEXVAL, ENFORCE, or REPORT phases, and the compatibility/concurrency considerations are different.

Note that TRIGGERs will fire for SHRLEVEL CHANGE but not for SHRLEVEL NONE.

REPLACE
Tells whether the table space and all its indexes need to be reset to empty

before records are loaded. With this option, the newly loaded rows replace *all* existing rows of all tables in the table space, not just those of the table you are loading. For DB2 STOGROUP-defined data sets, the data set is deleted and redefined with this option, unless you also specified the REUSE option. You must have LOAD authority for all tables in the table space where you perform LOAD REPLACE. If you attempt a LOAD REPLACE without this authority, you get an error message.

You cannot use REPLACE with the PART *integer* REPLACE option of INTO TABLE; you must either replace an entire table space using the REPLACE option or replace a single partition using the PART *integer* REPLACE option of INTO TABLE.

Specifying LOAD REPLACE (rather than PART *integer* REPLACE) tells LOAD to serialize at the table space level. If you want to serialize at the partition level, specify PART *integer* REPLACE. See the description for specifying REPLACE at the partition level under the keyword descriptions for INTO TABLE.

KEEPDICTIONARY

Prevents the LOAD utility from building a new compression dictionary. LOAD retains the current compression dictionary and uses it for compressing the input data. This option eliminates the cost associated with building a new dictionary.This keyword is valid only if the table space that is being loaded has the COMPRESS YES attribute.

If the table space or partition is empty, DB2 performs one of these actions:

- DB2 builds a dictionary if a compression dictionary does not exist.
- DB2 keeps the dictionary if a compression dictionary exists

If RESUME NO and REPLACE are specified when the table space or partition is not empty, DB2 performs the same actions as it does when the table space or partition is empty.

If the table space or partition is not empty and RESUME YES is specified, DB2 performs one of these actions:

- DB2 will not build a dictionary if a compression dictionary does not exist.
- DB2 keeps the dictionary if a compression dictionary exists.

For information regarding ESA data compression, see Part 5 (Volume 2) of *DB2 Administration Guide*.

COPYDDN *ddname1,ddname2*

Specifies the DD statements for the primary (*ddname1*) and backup (*ddname2*) copy data sets for the image copy.

ddname is the DD name.

The **default** is **SYSCOPY** for the primary copy.

The COPYDDN keyword can only be specified with REPLACE. A full image copy data set (SHRLEVEL REFERENCE) is created for the table or partitions specified when LOAD executes. The table space or partitions for which an image copy is produced is not placed in COPY-pending status.

Image copies taken during LOAD REPLACE are not recommended for use with RECOVER TOCOPY, because these image copies might contain unique index violations or referential constraint violations.

Using COPYDDN when loading a table with LOB columns will not create a copy of any index or LOB table space. You must perform these tasks separately.

The **COPYDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

RECOVERYDDN *ddname3,ddname4*
Specifies the DD statements for the primary (*ddname3*) and backup (*ddname4*) copy data sets for the image copy at the recovery site.

ddname is the DD name.

You cannot have duplicate image copy data sets. The same rules apply for RECOVERYDDN and COPYDDN.

The **RECOVERYDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

STATISTICS
Specifies the gathering of statistics for a table space, index, or both; the statistics are stored in the DB2 catalog.

If you specify the STATISTICS keyword with no other **statistics-spec** or **correlation-stats-spec** options, DB2 gathers only table space statistics. Statistics are collected on a base table space, but not on a LOB table space.

TABLE
Specifies the table for which column information is to be gathered. All tables must belong to the table space specified in the TABLESPACE option.

(ALL) Specifies that information is to be gathered for all columns of all tables in the table space.

SAMPLE *integer*
Indicates the percentage of rows to sample when collecting non-indexed column statistics. Any value from 1 through 100 can be specified. The **default** is **25**.

(*table-name***)**
Specifies the tables for which column information is to be gathered. The parentheses are required. If you omit the qualifier, the user identifier for the utility job is used.

If you specify more than one table, you must repeat the TABLE option.

COLUMN
Specifies columns for which column information is to be gathered.

You can only specify this option if you specify a particular tables for which statistics are to be gathered (TABLE (table-name)). If you specify particular tables and do not specify the COLUMN option, the default, COLUMN(ALL), is used. If you do not specify a

particular table when using the TABLE option, you cannot specify the COLUMN option; however, COLUMN(ALL) is assumed.

(ALL) Specifies that statistics are to be gathered for all columns in the table. The **default** is **ALL**.

(*column-name***, ...)**
Specifies the columns for which statistics are to be gathered. The parentheses are required.

You can specify a list of column names; the maximum is 10. If you specify more than one column, separate each name with a comma.

INDEX
Specifies indexes for which information is to be gathered. Column information is gathered for the first column of the index. All the indexes must be associated with the *same* table space, which must be the table space specified in the TABLESPACE option.

(ALL) Specifies that the column information is to be gathered for all indexes defined on tables contained in the table space. The parentheses are required.

(*index-name***)**
Specifies the indexes for which information is to be gathered. The parentheses are required.

REPORT
Determines if a set of messages is generated to report the collected statistics.

NO
Indicates that the set of messages is not output to SYSPRINT.

The **default** is **NO**.

YES
Indicates that the set of messages is output to SYSPRINT. The messages generated are dependent on the combination of keywords (such as TABLESPACE, INDEX, TABLE, and COLUMN) specified with the RUNSTATS utility. However, these messages are *not* dependent on the specification of the UPDATE option. REPORT YES always generates a report of SPACE and ACCESSPATH statistics.

UPDATE
Tells whether the collected statistics are inserted into the catalog tables. UPDATE also allows you to select statistics used for access path selection or statistics used by database administrators. What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

ALL Indicates that all collected statistics will be updated in the catalog.

The **default** is **ALL**.

ACCESSPATH
Indicates that only the catalog table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only the catalog table columns that provide statistics to help the database administrator assess the status of a particular table space or index are updated.

NONE Indicates that no catalog tables are updated with the collected statistics. This option is only valid when REPORT YES is specified.

HISTORY
Records all catalog table inserts or updates to the catalog history tables. What you specify under the UPDATE option affects what options you can specify under the HISTORY option. The table below lists the HISTORY options allowed for each UPDATE option specified.

The default will be supplied by the value specified in STATISTICS HISTORY on panel DSNTIPO.

ALL Indicates that all collected statistics will be updated in the catalog history tables.

ACCESSPATH
Indicates that only the catalog history table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only SPACE-related catalog statistics are updated in catalog history tables.

NONE Indicates that no catalog history tables are updated with the collected statistics.

Table 25. Allowable HISTORY options based on selected UPDATE option

UPDATE option	HISTORY options
ALL	ALL, ACCESSPATH, SPACE, NONE
ACCESSPATH	ACCESSPATH, NONE
SPACE	SPACE, NONE
NONE	NONE

FORCEROLLUP
Determines if aggregation or rollup of statistics will take place when RUNSTATS is executed in spite of some parts being empty. This will enable the optimizer to select the best access path.

The following are the available options for the FORCEROLLUP Keyword.

YES Indicates that forced aggregation or rollup processing is to be done, even though some parts may not contain data.

NO Indicates that aggregation or rollup will be done only if data is available for all parts.

If data is not available for all parts, DSNU623I message will be issued, provided the installation value for STATISTICS ROLLUP on panel DSNTIPO is set to NO.

KEYCARD
Collects all of the distinct values in all of the 1 to *n* key column combinations for the specified indexes. *n* is the number of columns in the index.

FREQVAL
Controls the collection of frequent value statistics. If you specify FREQVAL, it must be followed by two additional keywords:

NUMCOLS
Indicates the number of key columns to concatenate together when collecting frequent values from the specified index. Specifying '3' means to collect

frequent values on the concatenation of the first three key columns. The **default** is **1**, which means collect frequent values on the first key column of the index.

COUNT
Indicates the number of frequent values to be collected. Specifying '15' means collect 15 frequent values from the specified key columns. The **default** is **10**.

REUSE
When used with the REPLACE option, specifies that LOAD should logically reset and reuse DB2-managed data sets without deleting and redefining them. If you do not specify REUSE, DB2 deletes and redefines DB2-managed data sets to reset them.

REUSE must be accompanied by REPLACE to do the logical reset for all data sets. However, if you specify REUSE for the table space and REPLACE only at the partition level, only the replaced partitions are logically reset. See the description of REUSE in "Into-table-spec" on page 167 for information about specifying REUSE for individual partitions.

If a data set has multiple extents, the extents will not be released if you specify the REUSE parameter.

LOG Tells whether logging is to occur during the RELOAD phase of the load process.

YES
Specifies normal logging during the load process. All records loaded are logged.

The **default** is **YES**.

NO
Specifies no logging of data during the load process. The NO option sets the COPY-pending restriction against the table space or partition that the loaded table resides in. No table or partition in the table space can be updated until the restriction is removed. For ways to remove the restriction, see "Resetting the COPY-pending status" on page 209.

If you load a single partition of a partitioned table space and the table space has a nonpartitioning index, some logging might occur in the build phase as DB2 logs index structure changes. This logging allows recoverability of the nonpartitioning index in case an abend occurs, and also allows concurrency.

A LOB table space that was defined with LOG YES or LOG NO will affect logging while loading a LOB column. See Table 33 on page 204 for more information.

NOCOPYPEND
Specifies that LOAD is not to set the table space in the COPY-pending status, even though LOG NO was specified. A NOCOPYPEND specification will not turn on or change any informational COPY-pending (ICOPY) status for indexes. Normal completion of a LOAD LOG NO NOCOPYPEND job will be return code 0 if no other errors or warnings exist.

DB2 ignores a NOCOPYPEND specification if you also specified COPYDDN to make a local site inline image copy during the LOAD.

Attention: You should only specify the NOCOPYPEND option if the data in the table space can be easily recreated by another LOAD if it is lost. If you do not take an image copy following the LOAD, you will not be able to recover the table space using the RECOVER utility.

WORKDDN(*ddname1***,***ddname2***)**

Specifies the DD statements for the temporary work file for sort input and sort output. Temporary work files for sort input and output are required if the LOAD involves tables with indexes.

ddname1 is the DD name for the temporary work file for sort input.

The **default** is **SYSUT1**.

ddname2 is the DD name for the temporary work file for sort output.

The **default** is **SORTOUT**.

The **WORKDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

SORTKEYS *integer*

Specifies that index keys will be sorted in parallel during the SORTBLD phase to improve performance. Optionally, you may use *integer* to provide an estimate of the number of index keys to be sorted. The **default** is **0**.

SORTKEYS is recommended to improve performance unless the table space has no indexes, or it has only one index and the data being loaded is already sorted in key sequence. For more information about invoking a parallel index build, see "Improving performance with SORTKEYS" on page 196 and "Building indexes in parallel for LOAD" on page 200.

FORMAT

Identifies the format of the input record. If you use FORMAT, it uniquely determines the format of the input, and no field specifications are allowed in an INTO TABLE option. Follow FORMAT with either the UNLOAD or SQL/DS option.

If you omit FORMAT, the format of the input data is determined by the rules for field specifications described for the WHEN option of "Option descriptions for INTO TABLE" on page 169.

UNLOAD

Specifies that the input record format is compatible with the DB2 unload format. (The DB2 unload format is the result of REORG with the UNLOAD ONLY option.) Input records that were unloaded by the REORG utility are loaded into the tables from which they were unloaded, if there is an INTO TABLE option to specify each table. Do not add columns or change column definitions of tables between running REORG UNLOAD ONLY and LOAD FORMAT UNLOAD. Any WHEN clause on that statement is ignored; DB2 reloads the records into the same tables from which they were unloaded. This ensures that the input records are loaded into the proper tables. Input records that cannot be loaded are discarded. If the DCB RECFM parameter is specified on the DD statement for the input

data set, and the data set format has not been modified since the REORG UNLOAD (ONLY) operation, the record format must be variable (RECFM=V).

SQL/DS
Specifies that the input record format is compatible with the SQL/DS unload format. The data type of a column in the table to be loaded must be the same as the data type of the corresponding column in the SQL/DS table.

If the SQL/DS input contains rows for more than one table, the WHEN clause of the INTO TABLE option tells which input records load into which DB2 table.

For information on the correct DCB parameters to specify on the DD statement for the input data set, refer to *DB2 Server for VSE: DBS Utility* or *DB2 Server for VM: DBS Utility*.

LOAD cannot load SQL/DS strings that are longer than the DB2 limit. For information about DB2 limits, see "Appendix A. Limits in DB2 for OS/390 and z/OS" on page 627.

SQL/DS data that has been unloaded to disk under DB2 Server for VSE & VM resides in a simulated OS/390 type data set with a record format of VBS. That must be taken into consideration when transferring the data to another system to be loaded into a DB2 table (for example, the DB2 Server for VSE & VM FILEDEF must define it as an OS/390 type data set). Processing it as a standard CMS file puts the SQL/DS record type field at the wrong offset within the records; LOAD is unable to recognize them as valid SQL/DS input.

FLOAT
Specifies that LOAD is to expect the designated format for floating point numbers.

(S390) Specifies that LOAD is to expect that floating point numbers are provided in System/390® hexadecimal Floating Point (HFP) format. **(S390)** is the format that DB2 stores floating point numbers in. It is also the **default** if you do not explicitly specify the FLOAT keyword.

(IEEE) Specifies that LOAD is to expect that floating point numbers are provided in IEEE Binary Floating Point (BFP) format.

When you specify FLOAT(IEEE), DB2 converts the BFP data to HFP format as the data is being loaded into the DB2 table. If a conversion error occurs while converting from BFP to HFP, DB2 places the record in the discard file.

FLOAT(IEEE) is mutually-exclusive with any specification of the FORMAT keyword. If you specify both FLOAT(IEEE) and FORMAT, DB2 issues message DSNU070I.

BFP format is sometimes called IEEE Floating Point. For more information about the BFP format, see *ESA/390 Principles of Operation*.

EBCDIC
Specifies that the input data file is EBCDIC. The **default** is **EBCDIC**.

ASCII Specifies that the input data file is ASCII. Numeric, date, time, and timestamp internal formats are not affected by the ASCII option.

UNICODE
Specifies that the input data file is UNICODE. The UNICODE option does not affect the numeric, date, time, and timestamp formats.

CCSID
Specifies up to three coded character set identifiers (CCSIDs) for the input file. The first specifies the CCSID for SBCS data found in the input file, the second specifies the CCSID for mixed DBCS data, and the third specifies the CCSID for DBCS data. If any of these are specified as 0 or omitted, the CCSID of the corresponding data type in the input file is assumed to be the same as the installation default CCSID; that is, if the input data is EBCDIC, the omitted CCSIDs are assumed to be the EBCDIC CCSIDs specified at installation, and if the input data is ASCII, the omitted CCSIDs are assumed to be the ASCII CCSIDs specified at installation. If the CCSIDs of the input data file do not match the CCSIDs of the table being loaded, the input data is converted to the table CCSIDs before being loaded.

integer is any valid CCSID specification.

If the input data is UNICODE, the CCSIDs default to the UNICODE CCSIDs specified at system installation.

NOSUBS
Specifies that LOAD is not to accept substitution characters in a string.

A substitution character is sometimes placed in a string when that string is being converted from ASCII to EBCDIC, or converted from one CCSID to another. For example, this substitution occurs when a character (sometimes referred to as a code point) that exists in the source CCSID (code page) does not exist in the target CCSID (code page).

When you specify the NOSUBS option and the LOAD utility determines that a substitution character has been placed in a string as a result of a conversion, it performs one of the following actions:

- ***If discard processing is active:*** DB2 issues message DSNU310I and places the record in the discard file.
- ***If discard processing is not active:*** DB2 issues message DSNU334I and the utility abnormally terminates.

ENFORCE
Specifies whether or not LOAD is to enforce check constraints and referential constraints.

CONSTRAINTS
Indicates that constraints are to be enforced. If LOAD detects a violation, it deletes the errant row and issues a message to identify it. If you specify this option and referential constraints exist, sort input and sort output data sets are required.

The **default** is **CONSTRAINTS**.

NO Indicates that constraints are not to be enforced. This option places the target table space in the CHECK-pending status if at least one referential or check constraint is defined for the table.

ERRDDN *ddname*
Specifies the DD statement for a work data set for error processing. Information about errors encountered during processing is stored in this data set. A SYSERR data set is required if you request discard processing.

ddname is the DD name.

The **default** is **SYSERR**.

The **ERRDDN** keyword to specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

MAPDDN *ddname*
Specifies the DD statement for a work data set for error processing. It is used to map the identifier of a table row back to the input record that caused an error. A SYSMAP data set is *required* if you specify ENFORCE CONSTRAINTS and the tables have a referential relationship, or if you request discard processing when loading one or more tables that contain unique indexes.

ddname is the DD name.

The **default** is **SYSMAP**.

The **MAPDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

DISCARDDN *ddname*
Specifies the DD statement for a "discard data set", to hold copies of records that are not loaded (for example, if they contain conversion errors). The discard data set also holds copies of records loaded, then removed (due to unique index errors, or referential or check constraint violations). Input records can be flagged for discarding during RELOAD, INDEXVAL, and ENFORCE phases. However, the discard data set is not written until the DISCARD phase when the flagged records are copied from the input data set to the discard data set. The discard data set must be a sequential data set that can be written to by BSAM, with the same record format, record length, and block size as the input data set.

ddname is the DD name.

If you omit the DISCARDDN option, the utility application program saves discarded records only if there is a SYSDISC DD statement in the JCL input.

The **default** is **SYSDISC**.

The **DISCARDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses the DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

DISCARDS *integer*
Specifies the maximum number of source records to be written on the discard data set. *integer* can range from 0 to 2147483647. If the discard maximum is reached, LOAD abends, the discard data set is empty, and you cannot see which records were discarded. You can either restart the job with a larger limit, or terminate the utility.

DISCARDS 0 specifies that there is no maximum. The entire input data set can be discarded.

The **default** is **DISCARDS 0**.

SORTDEVT *device-type*
Specifies the device type for temporary data sets to be dynamically allocated by DFSORT. It can be any device type acceptable to the DYNALLOC parameter of the SORT or OPTION options for DFSORT.

If you omit SORTDEVT and a sort is required, you must provide the DD statements that the sort application program needs for the temporary data sets.

A TEMPLATE specification does not dynamically allocate sort work data sets. The SORTDEVT keyword controls dynamic allocation of these data sets.

SORTNUM *integer*
Tells the number of temporary data sets to be dynamically allocated by the sort application program.

If you omit SORTDEVT, SORTNUM is ignored. If you use SORTDEVT and omit SORTNUM, no value is passed to DFSORT. It is allowed to take its own default.

CONTINUEIF
Allows you to treat each input record as a portion of a larger record. After CONTINUEIF, write a condition in one of these forms:
(*start*:*end*) = X'*byte-string*'
(*start*:*end*) = '*character-string*'

If the condition is true in any record, the next record is concatenated with it before loading takes place. You can concatenate any number of records into a larger record up to a maximum size of 32767 bytes.

Data in the input record can be in American National Standard Code for Information Interchange (ASCII) or UNICODE, but the utility control statement always interprets character constants as EBCDIC. To use CONTINUEIF with the ASCII or UNICODE option, you must code the condition by using the hexadecimal form, not the character-string form. For example, use `(1:1)=X'31'` rather than `(1:1)='1'`.

(*start:end***)**
Are column numbers in the input record; the first column of the record is column 1. The two numbers tell the starting and ending columns of a continuation field in the input record.

Other field position specifications (such as those for WHEN, POSITION, or NULLIF) refer to the field position within the final assembled load record, not the input record.

The continuation field is removed from the input record and is not part of the final load record.

If you omit *:end*, the length of the continuation field is taken as the length of the byte string or character string. If you use *:end*, and the length of the resulting continuation field is not the same as the length of the byte string or character string, the shorter one is padded. Character strings are padded with blanks. Hexadecimal strings are padded with zeros.

X'*byte-string*'
Is a string of hexadecimal characters. That value in the continuation field indicates that the next input record is a continuation of the current load record. Records with that value are concatenated until the value in the continuation field changes. For example, a specification might be

```
CONTINUEIF (72) = X'FF'
```

'*character-string*'
Is a string of characters that has the same effect as X'*byte-string*'. For example, a specification might be

```
CONTINUEIF (99:100) = 'CC'
```

Into-table-spec

More than one table or partition for each table space can be loaded with a single invocation of the LOAD utility. At least one INTO TABLE statement is required for each table to be loaded, to:

- Identify the table that is to be loaded
- Describe fields within the input record
- Define the format of the input data set.

All tables specified by INTO TABLE statements must belong to the same table space.

If the data is already in UNLOAD or SQL/DS format, and FORMAT UNLOAD or FORMAT SQL/DS is used on the LOAD statement, no field specifications are allowed.

into-table-spec:

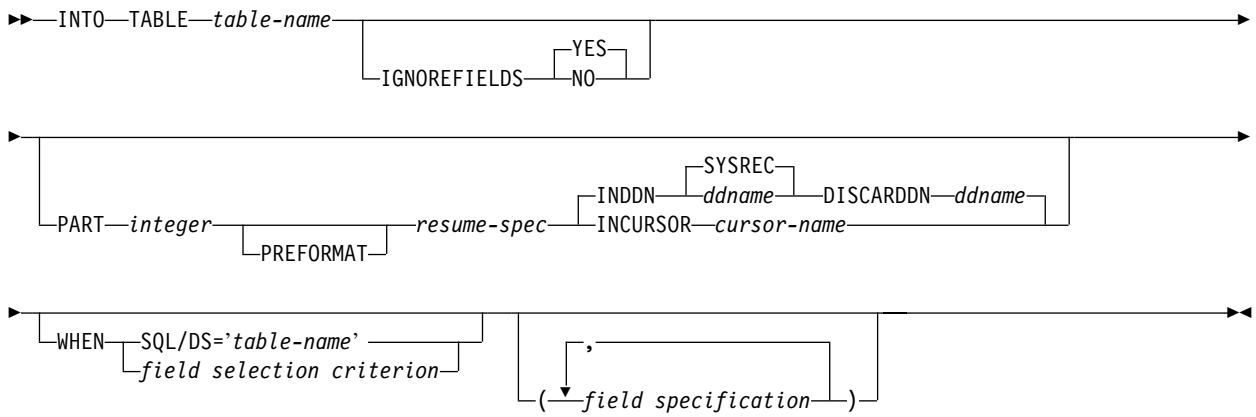

resume-spec:

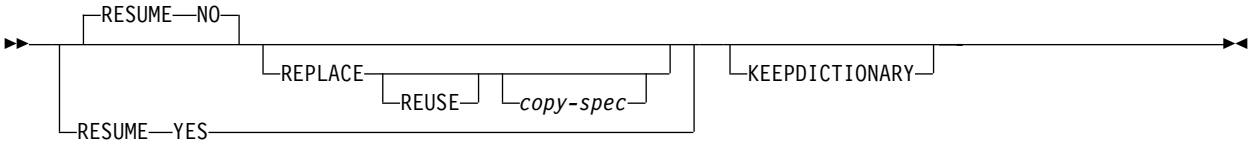

field selection criterion:

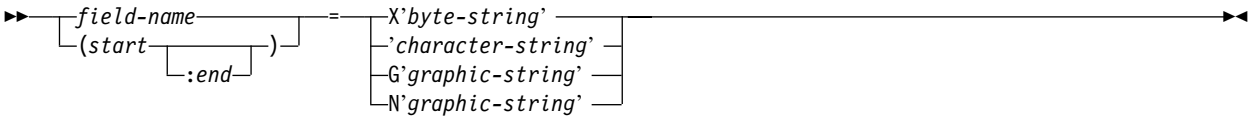

field specification:

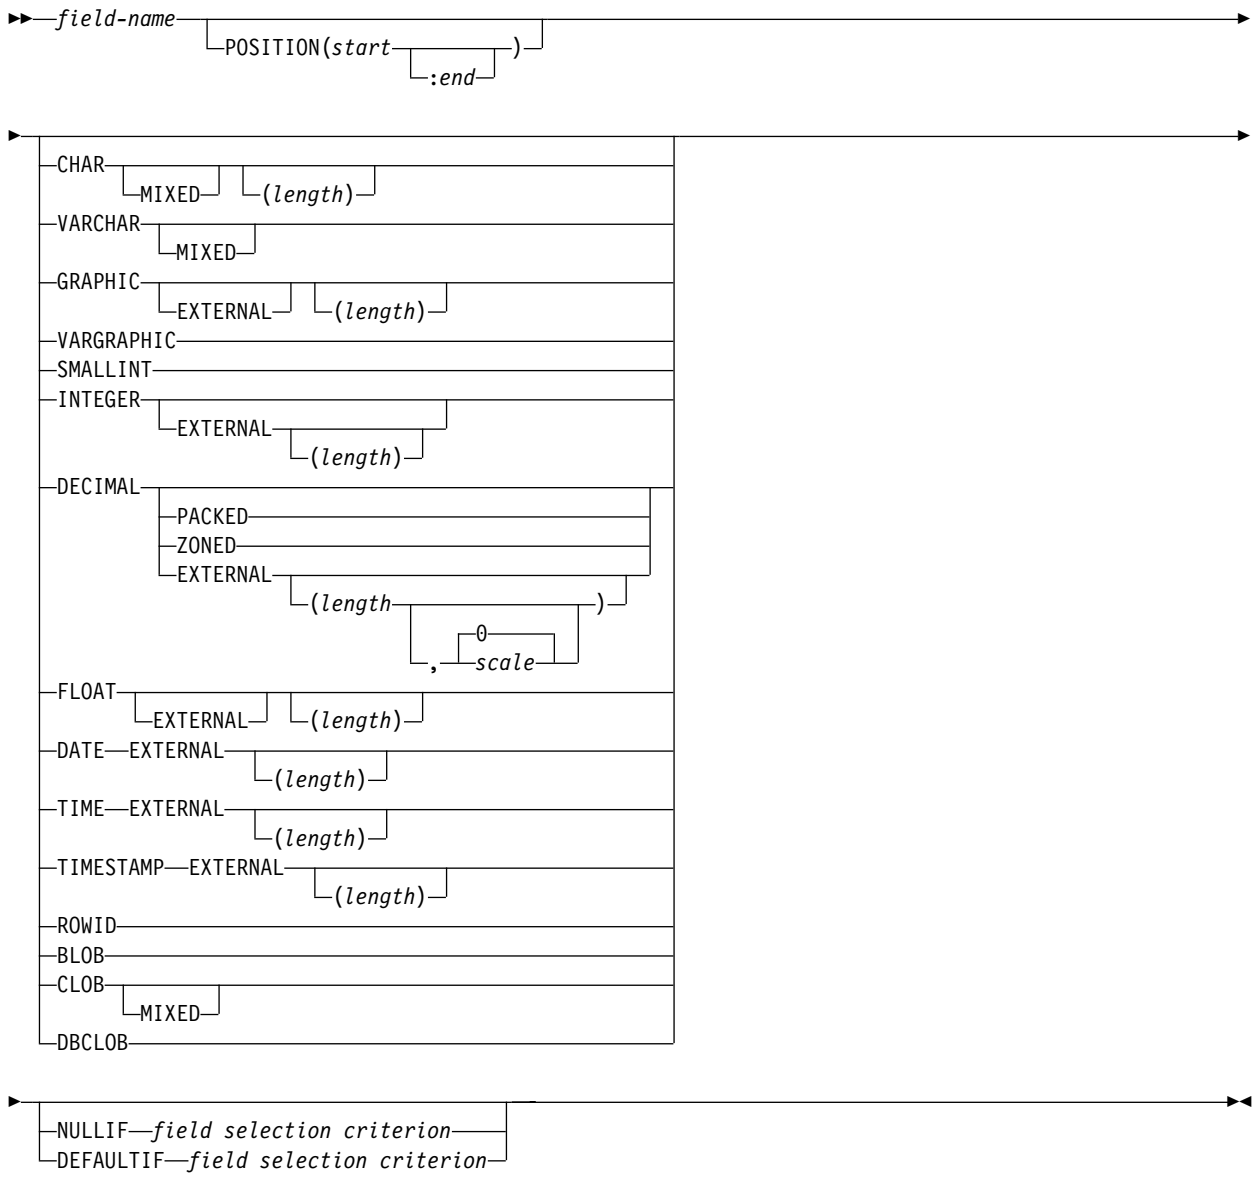

Option descriptions for INTO TABLE

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

table-name

Is the name of a table to be loaded. The table must be described in the catalog and must not be a catalog table.

If the table name is not qualified by an authorization ID, the authorization ID of the invoker of the utility job step is used as the qualifier of the table name.

Data from every load record in the data set is loaded into the table specified unless:

- A WHEN clause is used and the data does not match the field selection criterion
- The FORMAT UNLOAD option is used on the LOAD statement and the data comes from a table not specified in an INTO TABLE statement
- A certain partition is specified, and the data does not belong to that partition
- Data conversion errors occur
- Any errors not generated by data conversion occur

The following keywords are optional:

IGNOREFIELDS
Tells whether LOAD should skip fields in the input data set that do not correspond to columns in the target table. Examples of fields that do not correspond to table columns are the DSN_NULL_IND_nnnnn, DSN_ROWID, and DSN_IDENTITY fields that are generated by the REORG utility.

YES
Specifies that LOAD skips fields in the input data set that do not correspond to columns in the target table.

This can be useful if each input record contains a variable length field, followed by some variable length data that you do not wish to load and then some data that you want to load. Because of the variable length field, you cannot use the POSITION keyword to skip over the variable length data that you do not want to load. By specifying IGNOREFIELDS, you can give a field specification for the variable length data that you do not want to load; and by giving it a name that is not one of the table column names, it will be skipped over and not loaded.

This option should be used with care, since it will also cause fields to be skipped if you intend to load a column but have misspelled the name.

When you run the UNLOAD utility or the REORG utility with the UNLOAD EXTERNAL or DISCARD option, DB2 generates a LOAD statement that you can use to load the unloaded data into any table that has a compatible format. If the source table has a ROWID column that is defined with GENERATED ALWAYS, the generated LOAD statement contains a dummy field named DSN_IDENTITY for the identity column. The combination of IGNOREFIELDS and the dummy fields lets you load the unloaded data into a compatible table that has GENERATED ALWAYS columns.

If you want to include the data from the identity column or ROWID column when you load the unloaded data into a table, the identity column or ROWID column in the target table must be defined as GENERATED BY DEFAULT. To use the generated LOAD statement, remove the IGNOREFIELDS keyword and change the dummy field names to the names of the corresponding columns in the target table.

No
Specifies no skipping of fields in the load process.

PART *integer*

Is valid only for partitioned table spaces.

integer is the number of the partition for which records are accepted for loading. The same partition number cannot be specified more than once if partition parallelism has been requested. Any data outside the range of the specified partition is not loaded. The maximum is 254.

LOAD INTO PART *integer* is not allowed if an identity column is part of the partitioning index.

PREFORMAT

Specifies that the remaining pages are preformatted up to the high allocated RBA in the partition and its corresponding partitioning index space. The preformatting occurs after the data has been loaded and the indexes are built.

RESUME

Specifies whether records are to be loaded into an empty or non-empty partition. For nonsegmented table spaces, space occupied by rows that have been marked as deleted or by rows of dropped tables is not reused. If the RESUME option is specified at the table space level, the RESUME option is not allowed in the PART clause.

If you want the RESUME option to apply to the entire table space, use the LOAD RESUME option. If you want the RESUME option to apply to a particular partition, specify it using PART *integer* RESUME.

NO

Loads records into an empty partition. If the partition is not empty, and you have not used REPLACE, a message is issued, and the utility job step terminates with a job step condition code of 8.

For nonsegmented table spaces containing deleted rows or rows of dropped tables, using the REPLACE keyword provides increased efficiency.

The **default** is **NO**.

YES

Loads records into a non-empty partition. If the partition is empty, a warning message is issued, but the partition is loaded.

REPLACE

If specified, this option indicates that you want to replace only the contents of the partition cited by the PART option, rather than the entire table space.

You cannot use LOAD REPLACE with the PART *integer* REPLACE option of INTO TABLE. If you specify the REPLACE option, you must either replace an entire table space, using LOAD REPLACE, or a single partition, using the PART *integer* REPLACE option of INTO TABLE. You can, however, use PART *integer* REPLACE with LOAD RESUME YES.

REUSE

When used with the REPLACE option, specifies that LOAD should logically reset and reuse DB2-managed data sets without deleting and redefining them. If you do not specify REUSE, DB2 deletes and redefines DB2-managed data sets to reset them.

If you specify REUSE with REPLACE on the PART specification (and not for the “table space level” LOAD) only the specified partitions are logically

reset. If you specify REUSE for the table space and REPLACE for the partition, then data sets for the replaced parts are logically reset.

KEEPDICTIONARY

Prevents the LOAD utility from building a new dictionary. LOAD retains the current dictionary and uses it for compressing the input data. This option eliminates the cost associated with building a new dictionary.

This keyword is only valid if a dictionary exists and the partition being loaded has the COMPRESS YES attribute.

If the partition has the COMPRESS YES attribute, but there is no dictionary, one is built and an error message is issued.

INDDN *ddname*
Specifies the data definition (DD) statement or template identifying the input data set for the partition. The record format for the input data set must be fixed or variable. The data set must be readable by the basic sequential access method (BSAM).

The *ddname* is the name of the input data set.

The **default** is **SYSREC**.

If you specify INDDN, with or without DISCARDDN, in one INTO TABLE PART specification and you supply more than one INTO TABLE PART clause, you must specify INDDN in all INTO TABLE PART specifications.

Specifying INDDN at the part level and supplying multiple part clauses, each with their own INDDN, enables load partition parallelism, which can significantly improve performance. Loading all partitions in a single job with load partition parallelism is recommended instead of concurrent separate jobs whenever there are one or more NPIs on the table space.

The field specifications apply separately to each input file. Therefore, if multiple INTO TABLE PART INDDN clauses are used, field specifications are required on each one.

INDDN can be a template name.

DISCARDDN *ddname*
Specifies the DD statement for a discard data set for the partition. The discard data set holds copies of records that are not loaded (for example, if they contain conversion errors). The discard data set also holds copies of records that were loaded and then removed (due to unique index errors, or referential or check constraint violations).

Input records can be flagged for discarding during RELOAD, INDEXVAL, and ENFORCE phases. However, the utility does not write the discard data set until the DISCARD phase when the utility copies the flagged records from the input data set to the discard data set.

The discard data set must be a sequential data set, and it must be write-accessible by BSAM, with the same record format, record length, and block size as the input data set.

The *ddname* is the name of the discard data set.

If you omit the DISCARDDN option, LOAD does not save discarded records.

DISCARDDN can be a template name.

INCURSOR *cursor-name*

Specifies the cursor for the input data set. The cursor must be declared before it is used by the LOAD utility. You cannot load into the same table where you defined the cursor. Use the EXEC SQL utility control statement to define the cursor for the DB2 family cross loader function. The INCURSOR option is incompatible with the SHRLEVEL CHANGE option.

cursor-name is the cursor name.

The column names in the SELECT statement must be identical to the column names in the table being loaded. The AS clause in the select list can be used to change the columns names returned by the SELECT statement.

WHEN

The WHEN clause tells which records in the input data set are to be loaded. If there is no WHEN clause (and if FORMAT UNLOAD was not used in the LOAD statement), *all* records in the input data set are loaded into the specified tables or partitions. (Data beyond the range of the partition specified is not loaded.)

The option following WHEN describes a condition; input records that satisfy the condition are loaded. Input records that do not satisfy any WHEN clause of any INTO TABLE statement are written to the discard data set, if one is being used.

Data in the input record can be in ASCII or UNICODE, but the utility always interprets character constants that are specified in the utility control statement as EBCDIC. To use WHEN where the ASCII or UNICODE option is specified, code the condition using the hexadecimal form, not the character string form. For example, use (1:1)=X'31' rather than (1:1)='1'.

SQL/DS='*table-name*'

Is valid only when the FORMAT SQL/DS option is used on the LOAD statement.

table-name is the SQL/DS name of a table that has been unloaded onto the SQL/DS unload data set. The table name after INTO TABLE tells which DB2 table the SQL/DS table is loaded into.

If there is no WHEN clause, input records from every SQL/DS table are loaded into the table specified after INTO TABLE.

field-selection-criterion

Describes a field and a character constant. Only those records in which the field contains the specified constant are loaded into the table specified after INTO TABLE.

A field in a selection criterion must:

- Contain a character or graphic string. No data type conversions are performed when the contents of the field in the input record are compared to a string constant.
- Start at the same byte offset in each assembled input record. If any record contains varying-length strings—stored with length fields—that *precede* the selection field, then they must be padded so the start of the selection field is always at the same offset.

The field and the constant need not be the same length. If they are not, the shorter of the two is padded before a comparison is made. Character and graphic strings are padded with blanks. Hexadecimal strings are padded with zeros.

field-name
Is the name of a field defined by a *field-specification*. If *field-name* is used, the start and end positions of the field are given by the POSITION option of the field specification.

(start:end)
start and *:end* are column numbers in the assembled load record; the first column of the record is column 1. The two numbers tell the starting and ending columns of a selection field in the load record.

If *:end* is not used, the field is assumed to have the same length as the constant.

X'*byte-string***'**
Gives the constant as a string of hexadecimal characters. For example, write

```
WHEN (33:34) = X'FFFF'
```

'*character-string***'**
Gives the constant as a string of characters. For example, write

```
WHEN DEPTNO = 'D11'
```

G'*graphic-string***'**
Gives the constant as a string of double-byte characters. For example, write

```
WHEN (33:36) = G'<**>'
```

where “<” is the shift-out character, “*” is a double-byte character, and “>” is the shift-in character.

If the first or last byte of the input data is a shift-out character, it is ignored in the comparison. You can specify G either as an upper- or lower-case character.

N'*graphic-string***'**
Gives the constant as a string of double-byte characters. N and G are synonymous for specifying graphic string constants. You can specify N either as an upper- or lower-case character.

(*field-specification***, ...)**
Describes the location, format, and null value identifier of the data to be loaded.

If NO field specifications are used:

- The fields in the input records are assumed to be in the same order as in the DB2 table.
- The formats are set by the FORMAT option on the LOAD statement, if that is used.
- Fixed strings in the input are assumed to be of fixed maximum length. VARCHAR and VARGRAPHIC fields must contain a valid 2-byte binary length field preceding the data; there cannot be intervening gaps between them and the fields that follow.

- ROWID fields are varying length, and must contain a valid 2-byte binary length field preceding the data; there cannot be intervening gaps between them and the fields that follow.
- LOB data types are varying length, and require a valid 4-byte binary length field preceding the data; there cannot be intervening gaps between them and the fields that follow.
- Numeric data is assumed to be in the appropriate internal DB2 number representation.
- The NULLIF or DEFAULTIF options cannot be used.

If any field specification is used for an input table, there must be a field specification for each field of the table that does not have a default value. Any field in the table with no corresponding field specification is loaded with its default value.

If any column in the output table does not have a field specification and is defined as NOT NULL, with no default, the utility job step is terminated.

Identity columns may appear in the field specification only if they were defined with the GENERATED BY DEFAULT attribute.

field-name

Is the name of a field, and can be a name of your choice. If the field is to be loaded, the name must be the name of a column in the table named after INTO TABLE unless IGNOREFIELDS is specified. The field name can be used as a vehicle to specify the range of incoming data. See page 215.

The starting location of the field is given by the POSITION option. If POSITION is not used, the starting location is one column after the end of the previous field.

The length of the field is determined in one of the following ways, in the order listed:

1. If the field has data type VARCHAR, VARGRAPHIC, or ROWID, the length is assumed to be contained in a 2-byte binary field preceding the data. For VARCHAR fields, the length is in bytes; for VARGRAPHIC fields, it is in (double-byte) characters.

   If the field has data type CLOB, BLOB, or DBCLOB, the length is assumed to be contained in a 4-byte binary field preceding the data. For BLOB and CLOB fields, the length is in bytes; for DBCLOB fields, it is in (double-byte) characters.
2. If *:end* is used in the POSITION option, the length is calculated from *start* and *end*. In that case, any length attribute after the CHAR, GRAPHIC, INTEGER, DECIMAL, or FLOAT specifications is ignored.
3. The length attribute on the CHAR, GRAPHIC, INTEGER, DECIMAL, or FLOAT specifications is used as the length.
4. The length is taken from the DB2 field description in the table definition or assigned a default value according to the data type. For DATE and TIME fields the length is defined during installation. For variable length fields the length is defined from the column in the DB2 table definition, excluding the null indicator byte if it is present. Table 26 on page 176 shows the

default length, in bytes, for each data type.

Table 26. Default length of each data type (in bytes)

Data Type	Length in Bytes
BLOB	Varying
CHARACTER	Length used in column definition
CLOB	Varying
DATE	10 (or installation default)
DBCLOB	Varying
DECIMAL EXTERNAL	Same as for DECIMAL ZONED
DECIMAL PACKED	Length implied by column definition
DECIMAL ZONED	Decimal precision if decimal output column, otherwise length implied by column definition
FLOAT (single precision)	4
FLOAT (double precision)	8
GRAPHIC	2 * (length used in column definition)
INTEGER	4
MIXED	Mixed DBCS data
ROWID	Varying
SMALLINT	2
TIME	8 (or installation default)
TIMESTAMP	26
VARCHAR	Varying
VARGRAPHIC	Varying

If a data type is not given for a field, its data type is taken to be the same as that of the column it is loaded into, as given in the DB2 table definition.

POSITION(*start:end***)**

Tells where a field is in the assembled load record.

start and *end* are the locations of the first and last columns of the field; the first column of the record is column 1. The option can be omitted.

Column locations can be given as:

- An integer *n*, meaning an actual column number
- * , meaning one column after the end of the previous field
- *+*n*, where *n* is an integer, meaning *n* columns after the location specified by * .

The POSITION option specification *cannot* be enclosed in parentheses; however, the *start:end* description *must* be enclosed in parentheses, as the following example shows:

Valid	Invalid
`POSITION (10:20)`	`POSITION ((10:20))`

Data types in a field specification: The data type of the field can be specified by any of the keywords that follow. Except for graphic fields, *length* is the length in bytes of the input field.

All numbers designated EXTERNAL are in the same format in the input records.

CHAR

CHAR(*length***)**
For a fixed-length character string. The length of the string is determined from the POSITION specification or from *length*. You can also specify CHARACTER and CHARACTER(length).

MIXED
Specifies that the input field contains mixed SBCS and DBCS data. If MIXED is specified, then any required CCSID conversions use the mixed CCSID for the input data; if MIXED is not specified, then any such conversions will use the SBCS CCSID for the input data.

VARCHAR
For a character field of varying-length. The length in bytes must be given in a 2-byte binary field preceding the data. (The length given there does not include the 2 byte field itself.) The length field must start in the column specified as *start* in the POSITION option. If *:end* is used, it is ignored.

MIXED
Specifies that the input field contains mixed DBCS data. If MIXED is specified, then any required CCSID conversions use the mixed CCSID for the input data; if MIXED is not specified, then any such conversions will use the SBCS CCSID for the input data.

GRAPHIC EXTERNAL

GRAPHIC EXTERNAL(*length***)**
Used for a graphic field. You can specify both *start* and *end* for the field specification.

If you use GRAPHIC EXTERNAL, the input data must contain a shift-out character in the starting position, and a shift-in character in the ending position. Aside from the shift characters, there must be an even number of bytes in the field. The first byte of any pair must not be a shift character.

length is a number of double-byte characters. *length* for GRAPHIC EXTERNAL does not include the bytes of shift characters. The length of the field in bytes is twice the value of *length*.

For example, let *** represent 3 double-byte characters, and let < and > represent shift-out and shift-in characters. Then, to describe <***>, use either POS(1:8) GRAPHIC EXTERNAL or POS(1) GRAPHIC EXTERNAL(3).

GRAPHIC

GRAPHIC(*length***)**
Used for a graphic field. You can specify both *start* and *end* for the field specification.

If you use GRAPHIC, the input data must not contain shift characters. *start* and *end* must indicate the starting and ending positions of the data itself.

length is a number of double-byte characters. The length of the field in bytes is twice the value of *length*.

For example, let *** represent 3 double-byte characters. Then, to describe ***, use either POS(1:6) GRAPHIC or POS(1) GRAPHIC(3). A GRAPHIC field described in this way cannot be specified in a field selection criterion.

VARGRAPHIC
For a graphic field of varying-length. The length, in double-byte characters, must be given in a 2-byte binary field preceding the data. (The length given there does not include the 2 byte field itself.) The length field must start in the column specified as *start* in the POSITION option. *:end*, if used, is ignored.

VARGRAPHIC input data must not contain shift characters.

SMALLINT
For a two-byte binary number. Negative numbers are in two's complement notation.

INTEGER
Specifies a four-byte binary number. Negative numbers are in two's complement notation.

You can also specify INT.

INTEGER EXTERNAL

INTEGER EXTERNAL(*length***)**
A string of characters that represent a number. The format is that of an SQL numeric constant as described in Chapter 2 of *DB2 SQL Reference*.

You can also specify INT EXTERNAL.

DECIMAL

DECIMAL PACKED
For a number of the form *ddd...ds*, where *d* is a decimal digit represented by four bits, and *s* is a four-bit sign value. (The plus sign (+) is represented by A, C, E, or F and the minus sign (-) is represented by B or D.) The maximum number of *d*s is the same as the maximum number of digits allowed in the SQL definition.

You can also specify DEC or DEC PACKED.

DECIMAL ZONED
For a number of the form *znznzn...z/sn*, where *n* is a decimal digit represented by the right four bits of a byte (called the *numeric bits*); *z* is that digit's zone, represented by the left four bits; and *s* is the rightmost byte of the decimal operand, and can be treated as a zone or as the sign value for that digit. (The plus sign (+) is represented by A, C, E, or F and the minus sign (-) is represented by B or D.) The maximum number of *zn*s is the same as the maximum number of digits allowed in the SQL definition.

You can also specify DEC ZONED.

DECIMAL EXTERNAL

DECIMAL EXTERNAL(*length***)**

DECIMAL EXTERNAL(*length,scale***)**
A string of characters that represent a number. The format is that of an SQL numeric constant as described in Chapter 2 of *DB2 SQL Reference*.

length
Overall length of the input field in bytes.

scale

Specifies the number of digits to the right of the decimal point. That number must be an integer greater than or equal to 0, and can be greater than *length*.

The **default** is **0**.

If scale is greater than length, or the number of digits provided is less than the scale specified, the input number is padded on the left with zeros until the decimal point position is reached. If scale is greater than the target scale, the source scale locates the implied decimal position. All fractional digits greater than target scale are truncated. If scale is specified and the target column is small integer or integer, the decimal portion of the input number is ignored. If a decimal point is present, its position overrides the field specification of scale.

You can also specify DEC EXTERNAL and DEC EXTERNAL(*length*).

FLOAT EXTERNAL

FLOAT EXTERNAL(*length*)

A string of characters that represent a number. The format is that of an SQL floating point constant as described in Chapter 2 of *DB2 SQL Reference*.

If you specified the FLOAT(IEEE) or FLOAT(S390) option, it does not apply for this format (string of characters) of floating point numbers.

FLOAT(*length*)

For either a 64-bit floating point number, or a 32-bit floating point number. If *length* is between 1 and 21 inclusive, the number is 32 bits in the S390 (HFP) format:

Bit 0 Represents a sign (0 for ″plus″, and 1 for ″minus″)

Bits 1-7
Represent an exponent in excess-64 notation

Bits 8-31
Represent a mantissa.

If *length* is between 1 and 24 inclusive, the number is 32 bits in the IEEE (BFP) format:

Bit 0 Represents a sign (0 for ″plus″, and 1 for ″minus″)

Bits 1-8
Represent an exponent

Bits 9-31
Represent a mantissa.

If *length* is not specified, or is between 22 and 53 inclusive, the number is 64 bits in the S390 (HFP) format:

Bit 0 Represents a sign (0 for ″plus″, and 1 for ″minus″)

Bits 1-7
Represent an exponent in excess-64 notation

Bits 8-63
Represent a mantissa.

If *length* is not specified, or is between 25 and 53 inclusive, the number is 64 bits in the IEEE (BFP) format:

Bit 0 Represents a sign (0 for ″plus″, and 1 for ″minus″)

Bits 1-11
Represent an exponent in excess-64 notation

Bits 12-63
Represent a mantissa.

You can also specify REAL for single precision floating point and DOUBLE PRECISION for double precision floating point.

DATE EXTERNAL

DATE EXTERNAL(*length*)
For a character string representation of a date. Length, if unspecified, is the length given by the LOCAL DATA LENGTH install option, or, if none was provided, defaults to 10 bytes. If you specify a length, it must be within the range of 8 to 254 bytes.

Dates can be in any of the following formats. You can omit leading zeros for month and day. Trailing blanks can be included, but no leading blanks are allowed.
- *dd.mm.yyyy*
- *mm/dd/yyyy*
- *yyyy-mm-dd*
- any local format defined by your site at the time DB2 was installed.

TIME EXTERNAL

TIME EXTERNAL(*length*)
For a character string representation of a time. Length, if unspecified, is the length given by the LOCAL TIME LENGTH install option, or, if none was provided, defaults to eight bytes. If you specify a length, it must be within the range of 4 to 254 bytes.

Times can be in any of the following formats.
- *hh.mm.ss*
- *hh:mm* AM
- *hh:mm* PM
- *hh:mm:ss*
- any local format defined by your site at the time DB2 was installed.

You can omit the *mm* portion of the *hh:mm* AM and *hh:mm* PM formats if *mm* is equal to 00. For example, 5 PM is a valid time, and can be used instead of 5:00 PM

TIMESTAMP EXTERNAL

TIMESTAMP EXTERNAL(*length*)
For a character string representation of a time. *length*, if unspecified, defaults to 26 bytes. If you specify a length, it must be within the range of 19 to 26 bytes.

Timestamps can be in either of the following formats. Note that *nnnnnn* represents the number of microseconds, and can be from 0 to 6 digits. You can omit leading zeros from the month, day, or hour parts of the timestamp; you can omit trailing zeros from the microseconds part of the timestamp.
- *yyyy-mm-dd-hh.mm.ss*
- *yyyy-mm-dd-hh.mm.ss.nnnnnn*

See Chapter 2 of *DB2 SQL Reference* for more information about the timestamp data type.

ROWID
For a row ID. The input data must be a valid value for a row ID; DB2 will not perform any conversions.

A field specification for a row ID column is not allowed if the row ID column was created GENERATED ALWAYS.

If the ROWID column is part of the partitioning key, LOAD INTO TABLE PART is not allowed; specify LOAD INTO TABLE instead.

BLOB
For a BLOB field. You must specify the length in bytes in a 4-byte binary field preceding the data. (The length does *not* include the 4-byte field itself.) The length field must start in the column specified as *start* in the POSITION option. If *:end* is used, it is ignored.

CLOB
For a CLOB field. You must specify the length in bytes in a 4-byte binary field preceding the data. (The length does *not* include the 4-byte field itself.) The length field must start in the column specified as *start* in the POSITION option. If *:end* is used, it is ignored.

MIXED
Specifies that the input field contains mixed SBCS and DBCS data. If MIXED is specified, then any required CCSID conversions use the mixed CCSID for the input data; if MIXED is not specified, then any such conversions will use the SBCS CCSID for the input data.

DBCLOB
For a DBCLOB field. You must specify the length in double-byte characters in a 4-byte binary field preceding the data. (The length does *not* include the 4-byte field itself.) The length field must start in the column specified as *start* in the POSITION option. If *:end* is used, it is ignored.

Field selection criterion: The criterion describes a condition that causes the DB2 column to be loaded with NULL or its default value.

NULLIF *field-selection-criterion*
Describes a condition that causes the DB2 column to be loaded with NULL. The *field-selection-criterion* can be written with the same options as described on page 173. If the contents of the NULLIF field match the character constant given, the field specified in *field-specification* is loaded with NULL.

If the NULLIF field is defined by the name of a VARCHAR or VARGRAPHIC field, the length of the field is taken from the 2-byte binary field that appears before the data portion of the VARCHAR or VARGRAPHIC field.

| Data in the input record can be in ASCII or UNICODE, but the utility interprets character constants that are specified in the utility control statement always as EBCDIC. To use NULLIF when you specify the ASCII
| or UNICODE option, you must code the condition with the hexadecimal form, not the character string form. For example, use `(1:1)=X'31'` rather than `(1:1)='1'`.

The fact that a field in the output table is loaded with NULL does not change the format or function of the corresponding field in the input record. The input field can still be used in a field selection criterion. For example, with the field specification:

```
(FIELD1 POSITION(*) CHAR(4)
 FIELD2 POSITION(*) CHAR(3) NULLIF(FIELD1='SKIP')
 FIELD3 POSITION(*) CHAR(5))
```

and the source record:

```
SKIP   FLD03
```

the record is loaded so that:

FIELD1
Has the value 'SKIP'

FIELD2
Is NULL (not ' ' as in the source record)

FIELD3
Has the value 'FLD03'.

You cannot use the NULLIF parameter with the ROWID keyword, because row ID columns cannot be null.

DEFAULTIF *field-selection-criterion*
Describes a condition that causes the DB2 column to be loaded with its default value. The *field-selection-criterion* can be written with the same options as described on page 173. If the contents of the DEFAULTIF field match the character constant given, the field specified in *field-specification* is loaded with its default value.

If the DEFAULTIF field is defined by the name of a VARCHAR or VARGRAPHIC field, the length of the field is taken from the 2-byte binary field that appears before the data portion of the VARCHAR or VARGRAPHIC field.

Data in the input record can be in ASCII or UNICODE, but the utility interprets character constants that are specified in the utility control statement always as EBCDIC. When using DEFAULTIF when you specify the ASCII or UNICODE option, you must code the condition with the hexadecimal form, not the character string form. For example, use `(1:1)=X'31'` rather than `(1:1)='1'`.

You can use the DEFAULTIF attribute with the ROWID keyword. If the condition is met, the column will be loaded with a value generated by DB2.

Instructions for running LOAD

To run LOAD, you must:

1. Read “Before running LOAD” on page 183 in this chapter.
2. Prepare the necessary data sets, as described in “Data sets used by LOAD” on page 183.
3. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for LOAD, see “Sample control statements” on page 213.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 186. (For a complete description of the syntax and options for LOAD, see “Syntax and options of the control statement” on page 152.)
5. Check the compatibility table in “Concurrency and compatibility” on page 206 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the LOAD job doesn’t complete, as described in “Terminating or restarting LOAD” on page 204.
7. Read “After running LOAD” on page 208 in this chapter.
8. Run LOAD.

See “Chapter 3. Invoking DB2 online utilities” on page 23 for an explanation of ways to execute DB2 utilities.

Before running LOAD

Preprocessing input data: There is no sorting of the data rows during the LOAD utility— rows are loaded in the physical sequence in which they are found. It is a good idea to sort your input records in clustering sequence before loading.

You should also:

- Ensure that no duplicate keys exist for unique indexes.
- Correct check constraint violations and referential constraint violations in the input data set.

When loading into a segmented table space, sort your data by table to ensure that the data is loaded in the best physical organization.

Loading data using a cursor

Before you can load data using a cursor, you need to bind the DSNUTIL package at each location from which you plan to load data. Customize and run sample job DSNTIJXL to bind the packages. Sample job DSNTIJXL is in data set DSN710.SDSNSAMP.

Data sets used by LOAD

Table 27 describes the data sets used by LOAD. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 27. Data sets used by LOAD

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
Input data set	The input data set containing the data to be loaded. Its name or template is identified by the DD statement specified by the INDDN option. The default name is SYSREC. It must be a sequential data set that is readable by BSAM.	Yes
Sort data sets	Two temporary work data sets for sort input and sort output. Their DD or template names are specified with the WORKDDN option of the utility control statement. The default DD name for sort input is SYSUT1. The default DD name for sort output is SORTOUT.	Yes[1,3]
Mapping data set	Work data set for mapping the identifier of a table row back to the input record that caused an error. The default DD name is SYSMAP.	Yes[1,4]
UTPRINT	Contains messages from DFSORT (usually, SYSOUT or DUMMY).	No[5]

Table 27. Data sets used by LOAD (continued)

Data Set	Description	Required?
Discard data set	A work data set to hold copies of records not loaded. It must be a sequential data set that is readable by BSAM. Its DD or template name is specified with the DISCARDDN option of the utility control statement. If you omit this, LOAD creates the data set with the same record format, record length, and block size as the input data set. The default DD name is SYSDISC.	Yes[2]
Error data set	Work data set for error processing. Its DD name is specified with the ERRDDN parameter of the utility control statement. The default DD or template name is SYSERR.	Yes
Copy data sets	1 to 4 output data sets to contain image copy data sets. Their DD or template names are specified with the COPYDDN and RECOVERYDDN options of the utility control statement.	No

Note:
[1] When referential constraints exist and ENFORCE(CONSTRAINTS) is specified.
[2] If you request discard processing, by using the DISCARDS option of the utility control statement.
[3] For tables with indexes.
[4] If you request discard processing when loading one or more tables that have unique indexes.
[5] Required if a sort is done.

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table The name of the table to be loaded. It is named in the LOAD control statement and is accessed through the DB2 catalog. (If you want to load only one partition of a table, you must use the PART option in the control statement.)

Defining work data sets: Use Table 28 to calculate the size of work data sets for LOAD.

Table 28. Work data set calculation

Work Data Set	Size
SYSUT1	• Simple table space: *max(k,e)* • Partitioned or segmented table space: *max(k,e,m)* If SORTKEYS is used and an estimate of the number of keys to be sorted is specified: *max(f,e)* for a simple table space *max(f,e,m)* for a partitioned or segmented table space.
SORTOUT	*max(k,e)* If SORTKEYS is used, *max(f,e)*

Table 28. Work data set calculation (continued)

Work Data Set	Size
SYSERR	*e*
SYSMAP	• Simple table space or discard processing: *m* • Partitioned or segmented table space without discard processing: *max(m,e)*
SYSDISC	Same size as input data set

Notes:

1.
k = key calculation
f = foreign key calculation
m = map calculation
e = error calculation
max() = maximum value of the specified calculations

Calculating the key: k

max(longest index key + 14, longest foreign key + 14) × (number of keys extracted)

Calculating the number of keys extracted:

1. Count 1 for each index.
2. Count 1 for each foreign key that is not exactly indexed (that is, where foreign key and index definitions do not correspond identically).
3. For each foreign key that is exactly indexed (that is, where foreign key and index definitions correspond identically):
   a. Count 0 for the first relationship in which the foreign key participates.
   b. Count 1 for subsequent relationships in which the foreign key participates (if any).
4. Multiply count by the number of rows to be loaded.

Calculating the foreign key: f

max(longest foreign key + 14) × (number of keys extracted)

Calculating the map: m

The data set must be large enough to accommodate 1 map entry (length = 21 bytes) per table row produced by the LOAD job.

Calculating the error: e

The data set must be large enough to accommodate 1 error entry (length = 100 bytes) per defect detected by LOAD (for example, conversion errors, unique index violations, violations of referential constraints).

Calculating the number of possible defects:

– For discard processing, if the discard limit is specified, the number of possible defects is equal to the discard limit.

  If the discard limit is the maximum, the number of possible defects can be calculated as follows:

number of input records +
(number of unique indexes × number of keys extracted) +
(number of relationships × number of foreign keys extracted)

– For nondiscard processing, the data set is not required.

Allocating twice the space used by the input data sets is usually adequate for the sort work data sets. Two or three large SORTWK*nn* data sets are preferable to several small ones. For further information, see *DFSORT Application Programming: Guide.*

Instructions for specific tasks

The following tasks are described here:

"Loading variable-length data"
"Ordering loaded records"
"Replacing data with LOAD" on page 187
"Using LOAD for tables with identity columns or ROWID columns" on page 188
"Adding more data to a table or partition" on page 189
"Deleting all the data in a table space" on page 189
"Loading partitions" on page 189
"Loading data with referential constraints" on page 190
"Correcting referential constraint violations" on page 192
"Compressing data" on page 192
"Loading data from DL/I" on page 193
| "Loading data with a dynamic SQL statement" on page 194
"Using inline COPY with LOAD" on page 194
"Improving performance" on page 194
"Improving performance for parallel processing" on page 195
"Improving performance with SORTKEYS" on page 196
"Improving performance with LOAD or REORG PREFORMAT" on page 197

Loading variable-length data

To load variable-length data, put a 2-byte binary length field before each field of variable-length data. The value in that field depends on the data type of the column you load the data into. Use:

- The number of *single-byte characters* if the data type is VARCHAR
- The number of *double-byte characters* if the data type is VARGRAPHIC

For example, assume you have a variable-length column containing X'42C142C142C2', which might be interpreted as either six single-byte characters or three double-byte characters. With the two-byte length field, use:

- **X'0006'**X'42C142C142C2' for six single-byte characters in a VARCHAR column
- **X'0003'**X'42C142C142C2' for three double-byte characters in a VARGRAPHIC column

Ordering loaded records

LOAD loads records into a table space in the order in which they appear in the input stream. It does not sort the input stream, and does not insert records in sequence with existing records, even if there is a clustering index. To achieve clustering when loading an empty table or replacing data, sort the input stream. When adding data to a clustered table, consider reorganizing the table afterwards.

As rows with duplicate key values for unique indexes fail to be loaded, any records dependent on such rows will either:

- Fail to be loaded because they would cause referential integrity violations (if you specify ENFORCE CONSTRAINTS)
- Be loaded without regard to referential integrity violations (if you specify ENFORCE NO).

This might mean violations of referential integrity. Such violations can be detected by LOAD (without the ENFORCE(NO) option) or by CHECK DATA.

Replacing data with LOAD

You can use LOAD REPLACE to replace data in a single-table table space or in a multiple-table table space. You can replace all the data in a table space (using the REPLACE option), or you can load new records into a table space without destroying the rows already there (using the RESUME option).

If an object is in REORG-pending status, you can perform a LOAD REPLACE of the entire table space (which resets REORG-pending status). In this situation, no other LOAD operations are allowed. If a user defined table space is in refresh-pending (REFP) status, you can replace the data using LOAD REPLACE. See “Appendix C. Resetting an advisory or restrictive status” on page 677 for more information.

Using LOAD REPLACE with LOG YES: LOAD REPLACE or PART REPLACE with LOG YES logs only the reset and not each deleted row. If you need to see what rows are being deleted, use the SQL DELETE statement.

Replacing one table in a single-table table space: Figure 8 is an example that replaces one table in a single-table table space:

```
LOAD DATA
  REPLACE
  INTO TABLE DSN8710.DEPT
  ( DEPTNO    POSITION (1)     CHAR(3),
    DEPTNAME  POSITION (5)     VARCHAR,
    MGRNO     POSITION (37)    CHAR(6),
    ADMRDEPT  POSITION (44)    CHAR(3),
    LOCATION  POSITION (48)    CHAR(16) )
  ENFORCE NO
```

Figure 8. Example of using LOAD to replace one table in a single-table table space

Replacing one table in a multiple-table table space: When using LOAD REPLACE on a multiple-table table space, you must be careful, because LOAD works on an entire table space at a time. Thus, to replace all rows in a multiple-table table space, you have to work with one table at a time, using the RESUME YES option on all but the first table. For example, if you have two tables in a table space, you need to do the following:

1. Use LOAD REPLACE on the first table. This empties out the table space and replaces just the data for the first table.

```
LOAD DATA CONTINUEIF(72:72)='X'
  REPLACE
  INTO DSN8710.TOPTVAL
  ( MAJSYS    POSITION (2)    CHAR(1),
    ACTION    POSITION (4)    CHAR(1),
    OBJECT    POSITION (6)    CHAR(2),
    SRCHCRIT  POSITION (9)    CHAR(2),
    SCRTYPE   POSITION (12)   CHAR(1),
    HEADTXT   POSITION (80)   CHAR(50),
    SELTXT    POSITION (159)  CHAR(50),
    INFOTXT   POSITION (238)  CHAR(71),
    HELPTXT   POSITION (317)  CHAR(71),
    PFKTXT    POSITION (396)  CHAR(71),
    DSPINDEX  POSITION (475)  CHAR(2) )
```

Figure 9. Example of using LOAD REPLACE on the first table in a table space

2. Use LOAD with RESUME YES on the second table. This adds the records for the second table without destroying the data in the first table.

```
LOAD DATA CONTINUEIF(72:72)='X'
  RESUME YES
  INTO DSN8710.TDSPTXT
  ( DSPINDEX  POSITION (2)     CHAR(2),
    LINENO    POSITION (6)     CHAR(2),
    DSPLINE   POSITION (80)    CHAR(79) )
```

Figure 10. Example of using LOAD with RESUME YES on the second table in a table space

If you need to replace just one table in a multi-table table space, you need to delete all the rows in the table, then use LOAD with RESUME YES. For example, assume you want to replace all the data in DSN8710.TDSPTXT without changing any data in DSN8710.TOPTVAL. To do this:

1. Delete all the rows from DSN8710.TDSPTXT using an SQL DELETE statement. (The mass delete works most quickly on a segmented table space.)

   ```
   DELETE FROM DSN8710.TDSPTXT;
   ```

2. Use the LOAD job in Figure 11 to replace the rows in that table.

```
LOAD DATA CONTINUEIF(72:72)='X'
  RESUME YES
  INTO DSN8710.TDSPTXT
  ( DSPINDEX  POSITION (2)     CHAR(2),
    LINENO    POSITION (6)     CHAR(2),
    DSPLINE   POSITION (80)    CHAR(79) )
```

Figure 11. Example of using LOAD with RESUME YES to replace one table in a multi-table table space

Using LOAD for tables with identity columns or ROWID columns

When you run the UNLOAD utility or the REORG utility with the UNLOAD EXTERNAL or DISCARD options, DB2 generates a LOAD statement that you can use to load the unloaded data into any table that has a compatible format. If the source table has a ROWID column that is defined with GENERATED ALWAYS, the generated LOAD statement contains a dummy field named DSN_ROWID for the ROWID column. If the source table has an identity column that is defined with GENERATED ALWAYS, the generated LOAD statement contains a dummy field named DSN_IDENTITY for the identity column. The keyword IGNOREFIELDS in the LOAD statement causes DB2 to skip the DSN_ROWID or DSN_IDENTITY field

when it loads the data into a table. Using the combination of IGNOREFIELDS and the dummy fields, you can load the unloaded data into a compatible table that has GENERATED ALWAYS columns.

If you want to include the data from the identity column or ROWID column when you load the unloaded data into a table, the identity column or ROWID column in the target table must be defined as GENERATED BY DEFAULT. To use the generated LOAD statement, remove the IGNOREFIELDS keyword and change the dummy field names to the corresponding column names in the target table.

Adding more data to a table or partition

You may want to add data to a table, rather than replace it. The RESUME keyword specifies whether data is to be loaded into an empty or a non-empty table space. RESUME NO loads records into an empty table space. RESUME YES loads records into a non-empty table space.

If RESUME NO is specified and the target table is not empty, no data is loaded.

If RESUME YES is specified and the target table is empty, data IS loaded.

LOAD always adds rows to the end of the existing rows, but index entries are placed in key sequence.

Deleting all the data in a table space

Specifying LOAD REPLACE without loading any records is an efficient way of clearing a table space. To achieve this, the input data set should be specified in the JCL as DD DUMMY. LOAD REPLACE is efficient because:

1. LOAD REPLACE does not log any rows.
2. LOAD REPLACE redefines the table space.
3. LOAD REPLACE retains all views and privileges associated with a table space or table.
4. LOG YES can be used to make the LOAD REPLACE recoverable.

LOAD REPLACE will replace ALL TABLES in the table space.

Loading partitions

If you use the PART clause of the INTO TABLE option, only the specified partitions of a partitioned table are loaded. If you omit PART, the entire table is loaded.

You can specify the REPLACE and RESUME options separately by partition. The following example loads data into the first and second partitions of the employee table. Records with '0' in column 1 replace the contents of partition 1; records with '1' in column 1 are added to partition 2; all other records are ignored. (The example, simplified to illustrate the point, does not list field specifications for all columns of the table.)

Attention: If you are not loading columns in the same order as in the CREATE TABLE statement, you must code field specifications for each INTO TABLE statement.

```
LOAD DATA CONTINUEIF(72:72)='X'
  INTO TABLE DSN8710.EMP PART 1 REPLACE WHEN (1) = '0'
  ( EMPNO     POSITION (1:6)  CHAR(6),
    FIRSTNME  POSITION (7:18) CHAR(12),

   ⋮

                                          )
```

```
  INTO TABLE DSN8710.EMP PART 2 RESUME YES WHEN (1) = '1'
  ( EMPNO     POSITION (1:6)  CHAR(6),
    FIRSTNME  POSITION (7:18) CHAR(12),

  ⋮

                                        )
```

The following example assumes you have your data in separate input data sets. That data is already sorted by partition, so you do not have to use the WHEN clause of INTO TABLE. The RESUME YES option placed before the PART option inhibits concurrent partition processing while the utility is running.

```
LOAD DATA INDDN EMPLDS1 CONTINUEIF(72:72)='X'
  RESUME YES
  INTO TABLE DSN8710.EMP REPLACE PART 1

LOAD DATA INDDN EMPLDS2 CONTINUEIF(72:72)='X'
  RESUME YES
  INTO TABLE DSN8710.EMP REPLACE PART 2
```

The following example allows partitioning independence when loading more than one partition concurrently.

```
LOAD DATA INDDN SYSREC LOG NO
  INTO TABLE DSN8710.EMP PART 2 REPLACE
```

LOAD INTO PART *x* is not allowed if an identity column is part of the partitioning index.

| Specify a part clause with INDDN and optionally DISCARDDN keywords for each
| partition in your utility control statement to invoke partition parallelism. This will
| reduce the elapsed time required for loading large amounts of data into partitioned
| table spaces.

Consequences of DEFINE NO: If a partitioned table space is created with
DEFINE NO, all parts are DEFINE NO. The first data row inserted by the LOAD
utility defines all data sets in the partitioned table space. If this process takes a long
time, expect time-outs on the DBD. If you code your LOAD job with SHRLEVEL
CHANGE and use partition parallelism, it is equivalent to concurrent, independent
insert jobs. For example, in a large partitioned table space created with DEFINE
NO, the LOAD utility starts three tasks; the first task tries to insert the first row
causing an update to the DBD, the other two tasks time-out waiting to access the
DBD. The first task holds the lock on the DBD while the data sets get defined for
the table space.

Loading data with referential constraints

LOAD does not load a table with an incomplete definition; if the table has a primary key, then the unique index on that key must exist. If any table named to be loaded has an incomplete definition, the LOAD job terminates.

| LOAD requires access to the primary indexes on the parent tables of any loaded
| tables. For simple, segmented, and partitioned table spaces, it drains all writers
| from the parent table's primary indexes. Other users cannot make changes to the
| parent tables that result in an update to their own primary indexes. Concurrent
| inserts and deletes on the parent tables are blocked, but updates are allowed for
| columns that are not defined as part of the primary index.

By default, LOAD enforces referential constraints. By doing that, it provides you with several possibilities for error:

- Records to be loaded might have duplicate values of a primary key.
- Records to be loaded might have invalid foreign-key values, which are not values of the primary key of the corresponding parent table.
- The loaded table might lack primary key values that are values of foreign keys in dependent tables.

The next few sections describe how DB2 signals each of those errors and the means it provides for correcting them.

Duplicate values of a primary key: A primary index must be a unique index, and must exist if the table definition is complete. Therefore, when you load a parent table, you build at least its primary index. You need an error data set, and probably also a map data set and a discard data set.

Invalid foreign key values: A dependent table has the constraint that the values of its foreign keys must be values of the primary keys of corresponding parent tables. By default, LOAD enforces that constraint in much the same way as it enforces the uniqueness of key values in a unique index. First, it loads all records to the table; subsequently, it checks their validity with respect to the constraints, identifies any invalid record by an error message, and deletes the record. At your choice, the record can also be copied to a discard data set. Again you need at least an error data set, and probably also a map data set and a discard data set.

If a record fails to load because it violates a referential constraint, any of its dependent records in the same job also fail. For example, suppose that the sample project table and project activity tables belong to the same table space, that you load them both in the same job, and that some input record for the project table has an invalid department number. Then, that record fails to be loaded and does not appear in the loaded table; the summary report identifies it as causing a *primary* error.

But the project table has a primary key, the project number. In this case, the record rejected by LOAD defines a project number, and any record in the project activity table that refers to the rejected number is also rejected. The summary report identifies those as causing *secondary* errors. If you use a discard data set, both types of error records are copied to it.

Missing primary key values: The deletion of invalid records does not cascade to other dependent tables already in place. Suppose now that the project and project activity tables exist in separate table spaces, and that they are both currently populated and possess referential integrity. Further, suppose that the data in the project table is now to be replaced (using LOAD REPLACE) and that the replacement data for some department was inadvertently not supplied in the input data. Records referencing that department number might already exist in the project activity table. LOAD, therefore, automatically places the table space containing the project activity table (and all table spaces containing dependent tables of any table being replaced) into CHECK-pending status.

The CHECK-pending status indicates that the referential integrity of the table space is in doubt; it might contain records that violate a referential constraint. There are severe restrictions on the use of a table space in CHECK pending status; typically, you run the CHECK DATA utility to reset this status. For more information, see “Resetting the CHECK-pending status” on page 210.

Consequences of ENFORCE NO: If you use the ENFORCE NO option, you tell LOAD not to enforce referential constraints. Sometimes there are good reasons for

doing that but the result is that the loaded table space might violate the constraints. Hence, LOAD places the loaded table space in CHECK-pending status. If you use REPLACE, all table spaces containing any dependent tables of the tables that were loaded are also placed in CHECK-pending status. You must reset the status of each table before you can use any of the spaces.

Correcting referential constraint violations

The referential integrity checking in LOAD can only delete incorrect dependent rows, which were input to LOAD. Deletion is not always the best strategy for correcting referential integrity violations.

For example, the violations may occur because parent rows do not exist. In this case, it is better to correct the parent table, not to delete the dependent rows. Therefore and in this case, ENFORCE NO would be more appropriate than ENFORCE CONSTRAINTS. After the parent table is corrected, CHECK DATA can be used to reset the CHECK-pending status.

LOAD ENFORCE CONSTRAINTS is not equivalent to CHECK DATA. LOAD ENFORCE CONSTRAINTS deletes any rows causing referential constraint violations. CHECK DATA detects violations and optionally deletes such rows. CHECK DATA checks a complete referential structure, although LOAD checks only the rows being loaded.

When loading referential structures with ENFORCE CONSTRAINTS, parent tables should be loaded before dependent tables.

Compressing data

You can use LOAD with the REPLACE or RESUME NO options to build a *compression dictionary.* If your table space, or a partition in a partitioned table space, is defined with COMPRESS YES, the dictionary is created while records are loaded. After the dictionary is completely built, the rest of the data is compressed as it is loaded.

The data is not compressed until the dictionary is built. You must use LOAD REPLACE or RESUME NO to build the dictionary. To save processing costs, an initial LOAD does not go back to compress the records used to build the dictionary.

The number of records required to build a dictionary is dependent on the frequency of patterns in the data. For large data sets, the number of rows required to build the dictionary is a small percentage of the total number of rows to be compressed. For the best compression results, it is best to go ahead and build a new dictionary whenever you load the data.

However, there are circumstances in which you might want to compress data using an existing dictionary. If you are satisfied with the compression you are getting with an existing dictionary, you can keep that dictionary by using the KEEPDICTIONARY option of LOAD or REORG. For both LOAD and REORG, this method also saves you the processing overhead of building the dictionary.

Consider using KEEPDICTIONARY if the last dictionary was built by REORG; REORG's sampling method can yield more representative dictionaries than LOAD and can thus mean better compression. REORG with KEEPDICTIONARY is efficient because the data is not decompressed in the process.

However, REORG with KEEPDICTIONARY does not generate a compression report. You need to use RUNSTATS to update the catalog statistics and then query

the catalog columns yourself. See “Chapter 23. REORG TABLESPACE” on page 327 and “Chapter 26. RUNSTATS” on page 431 for more information about using REORG to compress data and about using RUNSTATS to update catalog information about compression.

Use KEEPDICTIONARY if you want to try to compress all the records during LOAD, and if you know the data has not changed much in content since the last dictionary was built. An example of LOAD with the KEEPDICTIONARY option is shown in Figure 12.

```
LOAD DATA
  REPLACE KEEPDICTIONARY
  INTO TABLE DSN8710.DEPT
  ( DEPTNO    POSITION (1)     CHAR(3),
    DEPTNAME  POSITION (5)     VARCHAR,
    MGRNO     POSITION (37)    CHAR(6),
    ADMRDEPT  POSITION (44)    CHAR(3),
    LOCATION  POSITION (48)    CHAR(16) )
  ENFORCE NO
```

Figure 12. Example of LOAD with the KEEPDICTIONARY option

You can also specify KEEPDICTIONARY for specific partitions of a partitioned table space. Each partition has its own dictionary.

Loading data from DL/I

To convert data in IMS DL/I databases from a hierarchic structure to a relational structure so that it can be loaded into DB2 tables, you can use the DataPropagator NonRelational (DPropNR) licensed program. DPropNR runs as an MVS application and can extract data from VSAM and physical sequential access method (SAM) files as well from DL/I databases.

Using DPropNR, you do not need to extract all the data in a database or data set. You use a statement such as an SQL subselect to tell which fields to extract and which conditions, if any, the source records or segments must meet.

With JCL models you edit, you can have DPropNR produce the statements for a DB2 LOAD utility job. If you have more than one DB2 system, you can name the one to receive the output. DPropNR can generate LOAD control statements in the job to relate fields in the extracted data to target columns in DB2 tables.

You can choose whether DPropNR writes the extracted data as either:

- 80-byte records included in the generated job stream
- A separate physical sequential data set, (which can be dynamically allocated by DPropNR) with a logical record length long enough to accommodate any row of the extracted data.

In the first case, the LOAD control statements generated by DPropNR include the CONTINUEIF option to describe the extracted data to DB2 LOAD.

In the second case, you can have DPropNR name the data set containing the extracted data in the SYSREC DD statement in the LOAD job. (In that case, DPropNR makes no provision for transmitting the extracted data across a network.)

Normally, you do not have to edit the job statements produced by DPropNR. However, in some cases you might have to edit; for example, if you want to load

character data into a DB2 column with INTEGER data type. (DB2 LOAD does not consider CHAR and INTEGER data compatible.)

DPropNR is a versatile tool that contains more control, formatting, and output options than are described here. For more information about them, see *DataPropagator NonRelational MVS/ESA Administration Guide.*

Loading data with a dynamic SQL statement

You can directly load the output of a SQL SELECT statement into a table on DB2 for OS/390. You can declare a cursor or execute one of the dynamic SQL statements listed in Table 19 on page 131 using the EXEC SQL statement. Use the result table from the cursor declared in the EXEC SQL statement as input to the LOAD utility statement defined with the INCURSOR option.

When the SQL statement is executed, the statement string is parsed and checked for errors. If the SQL statement is invalid, the error condition that prevented the execution is reported. If the SQL statement is valid but an error occurs during execution that error condition is reported. The utility terminates when it encounters an error.

Using inline COPY with LOAD

You can create a full image copy data set (SHRLEVEL REFERENCE) during LOAD execution. The new copy is an inline copy. The advantage to using inline copy is that the table space is not left in COPY-pending status regardless of which LOG option was specified for the utility. Thus, data availability is increased.

To create an inline copy, use the COPYDDN and RECOVERYDDN keywords. You can specify up to two primary and two secondary copies. Inline copies are produced during the RELOAD phase of LOAD processing.

The SYSCOPY record produced by an inline copy contains ICTYPE=F, SHRLEVEL=R. The STYPE column contains an R if the image copy was produced by LOAD REPLACE LOG(YES), and an S if the image copy was produced by LOAD REPLACE LOG(NO). The data set produced by the inline copy is logically equivalent to a full image copy with SHRLEVEL REFERENCE, but the data within the data set differs in some respects:

- Data pages might be out of sequence and some might be repeated. If pages are repeated, the last one is always the correct copy.
- Space map pages will be out of sequence and might be repeated
- If the compression dictionary is rebuilt with LOAD, the set of dictionary pages will occur twice in the data set, with the second set being the correct one.

The total number of duplicate pages will be small, with a negligible effect on the space required for the data set.

You must specify LOAD REPLACE. If you specify RESUME YES or RESUME NO but not REPLACE, an error message is issued and LOAD terminates.

Improving performance

To improve LOAD utility performance, you can:

- Use one LOAD DATA statement when loading multiple tables in the same table space. Follow the LOAD statement with multiple INTO TABLE WHEN statements.
- Run LOAD concurrently against separate partitions of a partitioned table space.

Alternatively, specify the INDDN and DISCARDDN keywords in your utility JCL to invoke partition parallelism. This will reduce the elapsed time required for loading large amounts of data into partitioned table spaces.

You should use load partition parallelism to load all partitions in a single job whenever there is one or more NPIs. If the only index is the partitioning index, then using multiple concurrent jobs against separate partitions is better.

- Preprocess the input data. For more information about preprocessing input data, see "Before running LOAD" on page 183.
- Load numeric data in its internal representation.
- Avoid data conversion, such as integer to decimal or decimal to floating-point.
- When you specify LOAD REPLACE, specify LOG NO with COPYDDN or RECOVERYDDN to create an inline copy.
- Sort the data in cluster order to avoid having to reorganize it after loading.
- Skip the sort phase of LOAD. The sort phase will be skipped when your input data meets all of the following conditions:
  - There is no more than one key per table.
  - All keys are of the same type (for example, all index keys, all indexed foreign keys, all nonindexed foreign keys, and so on).
  - The data being loaded is in key order.
  - The data being loaded is grouped by table and each input record must be loaded into one table only.
- If you cannot skip the sort phase because one or more of the conditions stated in the previous bullet are not met, use the SORTKEYS keyword to improve the efficiency of the sort.
- If you are using 3990 caching, and you have the nonpartitioning indexes on RAMAC®, consider specifying YES on the UTILITY CACHE OPTION field of installation panel DSNTIPE. This allows DB2 to use sequential prestaging when reading data from RAMAC for the following utilities:
  - LOAD PART integer RESUME
  - REORG TABLESPACE PART

  For these utilities, prefetch reads remain in the cache longer, thus possibly improving performance of subsequent writes.

The optimum order for presenting data to LOAD is as follows:

- If you are loading a single table that has, at most, one foreign key or one index key, sort the data in key sequence. (An index over a foreign key is allowed.) If it's an index key, sort the data in either ascending or descending order, depending on how the index was defined. If it's a foreign key, sort the data in ascending order. Null key values are treated as "high" values.
- If you are loading more than one table, choose one of the following methods:
  - Load each table separately. Using this method you can follow the rules listed above for loading single tables.
  - Use the WHEN clause under each INTO TABLE option on your LOAD statement to group your input data by table.

  Within each table, sort the data in key sequence.

Improving performance for parallel processing

Taking advantage of any new parallelism feature without allocating additional resources or tuning you system can lead to significant performance degradation. To

benefit from parallel operations when using online LOAD RESUME or parallel inserts, especially when non partitioning indexes are used, you can:

- Use a larger buffer pool to improve the buffer pool hit ratio.
- Define a higher deferred write threshold to reduce the number of pages written to disk which reduces the I/O time and contention.
- Define a larger checkpoint interval to reduce the number of pages written to disk which reduces the I/O time and contention.
- Use ESS Parallel Access Volume (PAV) to support multiple concurrent I/Os to the same volume containing non partitioning index data sets.
- Use non partitioning index pieces to support multiple concurrent non partitioning index I/Os.

Improving performance with SORTKEYS

Use the SORTKEYS keyword to improve performance of the index key sort.

Advantages of using SORTKEYS: If you use SORTKEYS, index keys are passed to sort in memory rather than written to work files. Avoiding this I/O to the work files improves LOAD performance.

You also reduce DASD space requirements for the SYSUT1 and SORTOUT data sets, especially if you provide an estimate of the number of keys to sort.

Using the SORTKEYS option reduces the elapsed time from the start of the reload phase to the end of the build phase.

However, if the index keys are already in sorted order, or there are no indexes, SORTKEYS does not provide any advantage.

You can reduce the elapsed time of a LOAD job for a table space or partition with more than one index defined by specifying the parameters to invoke a parallel index build. For more information, see “Building indexes in parallel for LOAD” on page 200.

Estimating the number of keys: You can specify an estimate of the number of keys for the job to sort. If the estimate is omitted or specified as 0, LOAD writes the extracted keys to the work data set, which reduces the performance improvement of using SORTKEYS.

To estimate the number of keys to sort:

1. Count 1 for each index
2. Count 1 for each foreign key where foreign key and index definitions are not identical
3. For each foreign key where foreign key and index definitions are identical:
   a. Count 0 for the first relationship in which the foreign key participates
   b. Count 1 for subsequent relationships in which the foreign key participates (if any).
4. Multiply the count by the number of rows to be loaded

If more than one table is being loaded, repeat the steps above for each table and sum the results.

Sort data sets: The space requirements for the sort input data set (SYSUT1) and the sort output data set (SORTOUT) are different if SORTKEYS is used. See page 184 for instructions on calculating the size of those data sets.

Improving performance with LOAD or REORG PREFORMAT

When DB2's preformatting delays impact the performance or execution time consistency of high INSERT applications and the table size can be predicted for a business processing cycle, LOAD PREFORMAT or REORG PREFORMAT might be a technique to consider. This technique will only be of value if DB2's preformatting causes a measurable delay with the INSERT processing or causes inconsistent elapsed times for INSERT applications. It is recommended that a performance assessment be conducted before and after LOAD or REORG PREFORMAT is used to quantify its value in your environment.

Considerations for using PREFORMAT: PREFORMAT is a technique used to eliminate DB2 having to preformat new pages in a table space during execution time. This might eliminate execution time delays but adds the preformatting cost as setup prior to the application's execution. LOAD or REORG PREFORMAT primes a new table space and prepares it for INSERT processing. When the preformatted space is utilized and DB2 has to extend the table space, normal data set extending and preformatting occurs.

Preformatting for INSERT processing may be desirable for high INSERT tables that will receive a predictable amount of data allowing all the required space to be pre-allocated prior to the application's execution. This would be the case for a table that acts as a repository for work items coming into a system that are subsequently used to feed a backend task that processes the work items.

Preformatting of a table space containing a table used for query processing may cause table space scans to read additional empty pages, extending the elapsed time for these queries. LOAD or REORG PREFORMAT is not recommended for tables that have a high ratio of reads to inserts if the reads result in table space scans.

Preformatting boundaries: You can manage your own data sets or have DB2 manage the data sets. For user-managed data sets, DB2 will not delete and reallocate them during utility processing. The size of the data set will not shrink back to the original data set allocation size but will either remain the same or increase in size if additional space or data is added. This has implications when LOAD or REORG PREFORMAT is used because preformatting causes all free pages between the high-used RBA (or page) to the high-allocated RBA to be preformatted. This includes secondary extents that may have been allocated.

For DB2 managed data sets, DB2 will delete and reallocate them if you specify REPLACE on the LOAD or REORG job. This will result in the data sets being re-sized to their original allocation size. They will remain that size if the data being reloaded does not fill the primary allocation and force a secondary allocation. This means the LOAD or REORG PREFORMAT option with DB2 managed data sets will at minimum cause the full primary allocation amount of a data set to be preformatted following the reload of data into the table space.

For both user-managed and DB2 managed data sets, if the data set goes into secondary extents during the utility processing, the high-allocated RBA becomes the end of the secondary extent and that becomes the high value for preformatting.

Preformatting performance considerations: LOAD or REORG PREFORMAT can eliminate dynamic preformatting delays when inserting into a new table space. The cost of this execution time improvement is an increase in the LOAD or REORG time due to the additional processing required to preformat all pages between the data loaded or reorganized and the high-allocated RBA. The additional LOAD or REORG time required depends on the amount of DASD space being preformatted.

Table space scans can also be elongated because empty preformatted pages will be read. It is best to use the LOAD or REORG PREFORMAT option for table spaces that start out empty and are filled through high insert activity before any query access is performed against the table space. Mixing inserts and nonindexed queries against a preformatted table space may impact the query performance without providing a compensating improvement in the insert performance. Best results may be seen where there is a high ratio of inserts to read operations.

Considerations for running LOAD

This section describes additional points to keep in mind when running LOAD.

Be aware that running the LOAD utility on a table space does not activate triggers defined on tables in the table space.

Converting input data

The LOAD utility converts data between compatible data types.[2]

Tables 29, 30, and 31identify the compatibility of data types for assignments and comparisons. Y indicates that the data types are compatible. N indicates no compatibility. D indicates the defaults used when you do not specify the input data type in a field specification of the INTO TABLE statement. For any number in a column, read the corresponding note at the bottom of the table.

Table 29. Numeric data conversion

Input Data Types	Output Data Types			
	SMALLINT	INTEGER	DECIMAL	FLOAT
SMALLINT	D	Y	Y	Y
INTEGER	Y	D	Y	Y
DECIMAL	Y	Y	D	Y
FLOAT	Y	Y	Y	D

Table 30. Character data conversion

Input Data Types	Output Data Types							
	BLOB	CHAR	VARCHAR	CLOB	GRAPHIC	VAR-GRAP	DBCLOB	ROWID
CHAR	Y	D	Y	Y	Y[1]	Y[1]	Y[1]	Y
CHAR MIXED	Y	D	Y	Y	Y[1]	Y[1]	Y[1]	N
VARCHAR	Y	Y	D	Y	Y[1]	Y[1]	Y[1]	Y

2. The source type is used for user-defined distinct types.

Table 30. Character data conversion (continued)

Input Data Types	Output Data Types							
	BLOB	CHAR	VARCHAR	CLOB	GRAPHIC	VAR-GRAP	DBCLOB	ROWID
VARCHAR MIXED	Y	Y	D	Y	Y[1]	Y[1]	Y[1]	N
GRAPHIC	N	Y[1]	Y[1]	Y[1]	D	Y	Y	N
VARGRAP	N	Y[1]	Y[1]	Y[1]	Y	D	Y	N
ROWID	N	N	N	N	N	N	N	D

Notes:

1. Conversion applies when either the input data or the target table is Unicode.

Table 31. Time data conversion

Input Data Types	Output Data Types		
	DATE	TIME	TIMESTAMP
DATE EXTERNAL	D	N	N
TIME EXTERNAL	N	D	N
TIMESTAMP EXTERNAL	Y	Y	D

Input fields with data types CHAR, CHAR MIXED, CLOB, DBCLOB, VARCHAR, VARCHAR MIXED, GRAPHIC, GRAPHIC EXTERNAL, and VARGRAPHIC are converted from the CCSIDs of the input file to the CCSIDs of the table space when they do not match. For example:

- You specify the American National Standard Code for Information Interchange (ASCII) or Unicode option for the input data and the table space is EBCDIC.
- You specify the EBCDIC or Unicode option and the table space is ASCII.
- You specify the ASCII or EBCDIC option and the table space is Unicode.
- The CCSID option is specified and the CCSIDs of the input data are not the same as the CCSIDs of the table space.

CLOB, BLOB, and DBCLOB input field types cannot be converted to any other field type.

Conversion errors cause LOAD:

- To abend, if there is no DISCARDS processing
- To map the input record for subsequent discarding and continue (if there is DISCARDS processing)

Truncation of the decimal part of numeric data is not considered a conversion error.

Specifying input fields

When specifying input fields, consider:

- Specify the length of VARCHAR, BLOB, CLOB, DBCLOB, and ROWID data in the input file.

- Explicitly define all input field specifications.
- Use DECIMAL EXTERNAL(length,scale) in full, or
- Specify decimal points explicitly in the input file.

Building indexes while loading data

LOAD builds all the indexes defined for any table being loaded. At the same time, it checks for duplicate values of any unique index key. If there are any duplicate values, none of the corresponding rows is loaded. Error messages identify the input records that produce duplicates; and, optionally, the records are copied to a discard data set. At the end of the job, a summary report lists all errors found.

For unique indexes, any two null values are taken to be equal, unless the index was created with the UNIQUE WHERE NOT NULL clause. In that case, if the key is a single column, it can contain any number of null values, though its other values must be unique.

Neither the loaded table nor its indexes contain any of the records that might have produced an error. Using the error messages, you can identify faulty input records, correct them, and load them again. If you use a discard data set, you can correct the records there and add them to the table with LOAD RESUME.

Building indexes in parallel for LOAD

Use parallel index build to reduce the elapsed time for a LOAD job by sorting the index keys and rebuilding multiple indexes in parallel, rather than sequentially. Optimally, a pair of subtasks process each index; one subtask sorts extracted keys while the other subtask builds the index. LOAD begins building each index as soon as the corresponding sort emits its first sorted record. For more information about improving index key sort performance, see “Improving performance with SORTKEYS” on page 196.

LOAD uses parallel index build if all of the following conditions are true:

- There is more than one index to be built.
- The LOAD job specifies the SORTKEYS keyword, *along with a non-zero estimate of the number of keys,* in the utility statement.
- You either allow the utility to dynamically allocate the data sets needed by SORT, or provide the necessary data sets yourself.

For a diagram of parallel index build processing, see Figure 20 on page 376.

Select one of the following methods to allocate sort work and message data sets:

Method 1: LOAD determines the optimal number of sort work and message data sets.

1. Specify the SORTKEYS and SORTDEVT keywords in the utility statement.
2. Allow dynamic allocation of sort work data sets by *not* supplying SORTWK*nn* DD statements in the LOAD utility JCL.
3. Allocate UTPRINT to SYSOUT.

Method 2: Allows you to control allocation of sort work data sets, while LOAD allocates message data sets.

1. Specify the SORTKEYS keyword in the utility statement.
2. Provide DD statements with DDNAMEs in the form SW*nn*WK*mm*.
3. Allocate UTPRINT to SYSOUT.

Method 3: Allows the most control over rebuild processing; you must specify both sort work and message data sets.

1. Specify the SORTKEYS keyword in the utility statement.
2. Provide DD statements with DDNAMEs in the form SW*nn*WK*mm*.
3. Provide DD statements with DDNAMEs in the form UTPRIN*nn*.

Data sets used: If you select Method 2 or 3 above, use the information provided here along with "Determining the number of sort subtasks", "Allocation of sort subtasks", and "Estimating the sort work file size" on page 202 to define the necessary data sets.

Each sort subtask must have its own group of sort work data sets and its own print message data set. Possible reasons to allocate data sets in the utility job JCL rather than using dynamic allocation are:

- To control the size and placement of the data sets.
- To minimize device contention.
- To optimally utilize DASD free space.
- To limit the number of utility subtasks used to build indexes.

The DDNAMEs SW*nn*WK*mm* define the sort work data sets used during utility processing. *nn* identifies the subtask pair, while *mm* identifies one or more data sets to be used by that subtask pair. For example:

SW01WK01 The first sort work data set used by the subtask building the first index.

SW01WK02 The second sort work data set used by the subtask building the first index.

SW02WK01 The first sort work data set used by the subtask building the second index.

SW02WK02 The second sort work data set used by the subtask building the second index.

The DDNAMEs UTPRIN*nn* define the sort work message data sets used by the utility subtask pairs. *nn* identifies the subtask pair.

Determining the number of sort subtasks: The maximum number of utility subtask pairs started for parallel index build is equal to the number of indexes to be built.

LOAD determines the number of subtask pairs according to the following guidelines:

- The number of subtask pairs equals the number of sort work data set groups allocated.
- The number of subtask pairs equals the number of message data sets allocated.
- If you allocate both sort work and message data set groups, the number of subtask pairs equals the smallest number of data sets allocated.

Allocation of sort subtasks: LOAD attempts to assign one sort subtask pair for each index to be built. If LOAD cannot start enough subtasks to build one index per subtask pair, it allocates any excess indexes across the pairs (in the order that the indexes were created), so one or more subtask pairs might build more than one index.

During parallel index build processing, LOAD assigns all foreign keys to the first utility subtask pair. Remaining indexes are then distributed among the remaining

subtask pairs according to the creation date of the index. If a table space does not participate in any relationships, LOAD distributes all indexes among the subtask pairs according to the index creation date, assigning the first created index to the first subtask pair.

Refer to Table 32 for conceptual information about subtask pairing when the number of indexes (seven indexes) exceed the available number of subtask pairs (five subtask pairs).

Table 32. LOAD subtask pairing for a relational table space

Subtask Pair	Index Assigned
SW01WK*mm*	Foreign keys, Fifth created index
SW02WK*mm*	First created index, Sixth created index
SW03WK*mm*	Second created index, Seventh created index
SW04WK*mm*	Third created index
SW05WK*mm*	Fourth created index

Estimating the sort work file size: If you choose to provide the data sets, you will need to know the size and number of keys present in all of the indexes being processed by the subtask in order to calculate each sort work file size. After you've determined which indexes are assigned to which subtask pairs, use the following formula to calculate the space required:

2 × (longest index key + 14) × (number of keys extracted)

longest key — The length of the longest key that will be processed by the subtask. For the first subtask pair for LOAD, compare the length of the longest key and the longest foreign key, and use the largest value.

number of keys — The number of keys from all indexes to be sorted that will be processed by the subtask.

Leaving free space

When loading into a nonsegmented table space, LOAD leaves one free page after reaching the FREEPAGE limit, regardless of whether the records loaded belong to the same or different tables.

When loading into a segmented table space, LOAD leaves free pages, and free space on each page, in accordance with the current values of the FREEPAGE and PCTFREE parameters. (Those values can be set by the CREATE TABLESPACE, ALTER TABLESPACE, CREATE INDEX, or ALTER INDEX statements.) LOAD leaves one free page after reaching the FREEPAGE limit for each table in the table space.

Loading with RECOVER-pending or REBUILD-pending status

You cannot load records specifying RESUME YES if any partition of a table space is in the RECOVER-pending status. In addition, you cannot load records if any index on the table being loaded is in the REBUILD-pending status. for information about resetting the RECOVER-pending status. See "Resetting the REBUILD-pending status" on page 269 for information about resetting the REBUILD-pending status.

If you are replacing a partition, these restrictions are relaxed; the partition being replaced can be in the RECOVER-pending status, and its corresponding index

partition can be in the REBUILD-pending status. However, all nonpartitioning indexes must *not* be in the page set REBUILD-pending status. See “Appendix C. Resetting an advisory or restrictive status” on page 677 for more information about resetting a restrictive status.

There is one RECOVER-pending restrictive status:

RECP The table space or partition is in the RECOVER-pending status. If a single logical partition is in RECP, the partition is treated as RECP for SQL access. A single logical partition in RECP does not restrict utility access to other logical partitions not in RECP. RECP is reset by recovering only the single logical partition.

There are three REBUILD-pending restrictive states:

RBDP REBUILD-pending status (RBDP) is set on a physical or logical index partition. The individual physical or logical partition is inaccessible and must be rebuilt using the REBUILD INDEX utility, or recovered using the RECOVER utility.

PSRBD

Page set REBUILD-pending (PSRBD) is set for nonpartitioning indexes. The entire index space is inaccessible and must be rebuilt using the REBUILD utility, or recovered using the RECOVER utility.

RBDP*

RBDP* (REBUILD-pending star) status is set only on logical partitions of nonpartitioning indexes. The entire index is inaccessible, but is made available again when the affected partitions are rebuilt using the REBUILD INDEX utility, or recovered using the RECOVER utility.

See Table 134 on page 681 for information about resetting the RECOVER pending status, and Table 133 on page 680 for information about resetting the REBUILD-pending status.

Using exit procedures

Any field procedure associated with a column of a table being loaded is executed to encode the data before it is loaded. The field procedures for all columns are executed before any edit or validation procedure for the row.

Any field specification that describes the data is checked before a field procedure is executed. That is, the field specification must describe the data as it appears in the input record.

Loading columns defined as ROWID

Columns defined as ROWID can be designated as input fields using the LOAD field specification syntax diagram. LOAD PART is not allowed if the ROWID column is part of the partitioning key. In this situation, DB2 issues error message DSNU256I.

Columns defined as ROWID can be designated as GENERATED BY DEFAULT or GENERATED ALWAYS. With GENERATED ALWAYS, DB2 always generates a Row ID.

ROWID generated by default: Columns defined as ROWID GENERATED BY DEFAULT can be set by the LOAD utility from input data. The input field must be specified as a ROWID. No conversions are allowed. The input data for a ROWID column must be a unique, valid value for a row ID. If the value of the row is not unique, a duplicate key violation will occur. If such an error occurs, the load will fail. In this case, you need to discard the duplicate value and retry the load with a new unique value, or allow DB2 to generate the value of the row ID.

You can use the DEFAULTIF attribute with the ROWID keyword. If the condition is met, the column will be loaded with a value generated by DB2. You cannot use the NULLIF attribute with the ROWID keyword, because row ID columns cannot be null.

ROWID generated always: A ROWID column that is defined as GENERATED ALWAYS cannot be included in the field specification list, because DB2 generates the row ID value for you.

Loading a LOB column

LOB columns are treated by the LOAD utility as varying-length data. The length value for a LOB column must be 4 bytes. When the input record is greater than 32KB, you might have to load the LOB data separately. See sample job DSN8DLPL in SDSNSAMP for an example.

Using LOAD LOG on a LOB table space

A LOB table space that was defined with LOG YES or LOG NO will affect logging while loading a LOB column. Table 33 shows the logging output and LOB table space effect, if any.

Table 33. LOAD LOG and REORG LOG impact for a LOB table space

LOAD LOG/ REORG LOG keyword	LOB table space LOG attribute	What is logged	LOB table space status after utility completes
LOG YES	LOG YES	Control information and LOB data	No pending status
LOG YES	LOG NO	Control information	No pending status
LOG NO	LOG YES	Nothing	COPY Pending
LOG NO	LOG NO	Nothing	COPY Pending

Inline statistics collection for discarded rows

If you specify the DISCARDDN and STATISTICS options and a row is found with check constraint errors or conversion errors, the row is not loaded into the table and DB2 does not collect inline statistics on it. However, LOAD utility processing collects inline statistics prior to discarding rows as a result of unique index violations or referential integrity violations. In these cases, if the number of discarded rows is large enough to make the statistics significantly inaccurate, run the RUNSTATS utility separately on the table to gather the most accurate statistics.

Inline COPY for a base table space

If you take an inline image copy of a table with LOB columns, DB2 makes a copy of the base table space, but does not copy the LOB table spaces.

Terminating or restarting LOAD

For instructions on restarting a utility job, see “Restarting an online utility” on page 46.

Terminating LOAD: If you terminate LOAD (LOAD SHRLEVEL CHANGE has different requirements) using the TERM UTILITY command during the reload phase, the records are not erased. The table space remains in RECOVER-pending status, and indexes remain in the REBUILD-pending status.

If you terminate LOAD using the TERM UTILITY command during the sort or build phases, then the indexes not yet built remain in the REBUILD pending status.

If you use the SORTKEYS option and the LOAD job terminates during the RELOAD, SORT, BUILD, or SORTBLD phases, then both RESTART and RESTART(PHASE) restart from the beginning of the RELOAD phase. However, restart of LOAD RESUME YES or LOAD PART RESUME YES in the BUILD or SORTBLD phase will result in message DSNU257I.

Table 34. LOAD phases and pending statuses

Phase	Effect on Pending Status
Reload	Places table space in RECOVER-pending status, then resets the status. Places indexes in REBUILD pending status. Places table space in COPY-pending status. Places table space in CHECK-pending status. Resets COPY-pending at end of phase if an inline copy is produced unless SORTKEYS is also specified.
Build	Resets REBUILD-pending status for nonunique indexes. Resets COPY-pending status at end of phase if an inline copy is produced and SORTKEYS is also specified.
Indexval	Resets REBUILD-pending status for unique indexes.
Enforce	Resets CHECK-pending status for table space.

Restarting LOAD: Table 35 provides information about restarting LOAD, depending on the phase LOAD was in when the job stopped.

- If you restart LOAD during the UTILINIT phase, it re-executes from the beginning of the phase.
- If LOAD abends or system failure occurs while it is in the UTILTERM phase, you must restart with RESTART(PHASE).
- If you restart a LOAD job for a table with LOB columns that specified the RESUME YES option, you must use RESTART CURRENT.
- If you use RESTART PHASE to restart a LOAD job which specified RESUME NO, the LOB table spaces and indexes on auxiliary tables will be reset.
- For a table with LOB columns, you cannot restart a LOAD job that specifies the INCURSOR option.
- If you restart a LOAD job which uses the STATISTICS keyword, inline statistics collection will not occur. To update catalog statistics, run the RUNSTATS utility after the restarted LOAD job completes.

In this table, the TYPE column distinguishes between the effects of specifying RESTART or RESTART(PHASE).

Table 35. LOAD restart information

Phase	Type	Data Sets Required	Notes
RELOAD	CURRENT	SYSREC and SYSUT1	
		SYSMAP and SYSERR	2,3
	PHASE	SYSREC	7
SORT	CURRENT	SYSUT1	1
	PHASE	SYSUT1	
BUILD	CURRENT	SORTOUT	1,5
	PHASE	SORTOUT	5
SORTBLD	CURRENT	SYSUT1 and SORTOUT	5,9
	PHASE	SYSUT1 and SORTOUT	5,9
INDEXVAL	CURRENT	SYSERR or SYSUT1	3
	PHASE	SYSERR or SYSUT1	3

Table 35. LOAD restart information (continued)

Phase	Type	Data Sets Required	Notes
ENFORCE	CURRENT	SORTOUT and SYSUT1	4
	PHASE	SORTOUT and SYSUT1	4
DISCARD	CURRENT	SYSMAP and SYSERR	4
		SORTOUT and SYSUT1	8
	PHASE	SYSMAP and SYSERR	4
		SORTOUT and SYSUT1	8
REPORT	CURRENT	SYSERR or SORTOUT	4
		SYSMAP and SYSERR	6
	PHASE	SYSERR or SORTOUT	4
		SYSMAP and SYSERR	6

Note:
1. The utility can be restarted with either RESTART or RESTART(PHASE). However, because this phase does not take checkpoints, RESTART is always re-executed from the beginning of the phase.
2. SYSMAP and SYSERR data sets may not be required for all load jobs. See "Chapter 14. LOAD" on page 151 for exact requirements.
3. If the SYSERR data set is not required and has not been provided, LOAD uses SYSUT1 as a work data set to contain error information.
4. This utility can be restarted with either RESTART or RESTART(PHASE). However, the utility can be re-executed from the last internal checkpoint. This is dependent on the data sets used and whether any input data sets have been rewritten.
5. LOAD RESUME YES cannot be restarted in the BUILD or SORTBLD phase.
6. If report is required and this is a load without discard processing, SYSMAP is required to complete the report phase.
7. You must not restart during RELOAD phase if you specified SYSREC DD *. This prevents internal commits from being taken, and RESTART performs like RESTART(PHASE), except with no data back-out. Also, you must not restart if your SYSREC input consists of multiple, concatenated data sets.
8. The SYSUT1 data set is required if the target table space is segmented or partitioned.
9. If you specified SORTKEYS, then use RESTART or RESTART(PHASE) to restart at the beginning of the RELOAD phase.

You can restart LOAD at its last commit point or at the beginning of the phase during which operation ceased. LOAD output messages identify the completed phases; use the DISPLAY command to identify the specific phase during which operation stopped.

Restarting after an out of space condition

See "Restarting after the output data set is full" on page 47 for guidance in restarting LOAD from the last commit point after receiving an out of space condition.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

For nonpartitioning indexes, LOAD PART:
- Drains only the logical partition
- Does not set the page set REBUILD-pending status (PSRBD)
- Does not respect PCTFREE or FREEPAGE attributes when inserting keys

Claims and drains

Table 36 shows which claim classes LOAD drains and the restrictive states the utility sets.

Table 36. Claim classes of LOAD operations. Use of claims and drains; restrictive states set on the target object.

Target	LOAD SHRLEVEL NONE	LOAD PART SHRLEVEL NONE	LOAD SHRLEVEL CHANGE	LOAD PART SHRLEVEL CHANGE
Table space, index, or physical partition of a table space or index	DA/UTUT	DA/UTUT	CW/UTRW	CW/UTRW
Nonpartitioning index	DA/UTUT	DR	CW/UTRW	CW/UTRW
Index logical partition		DA/UTUT		CW/UTRW
Primary index (with ENFORCE option only)	DW/UTRO	DW/UTRO	CR/UTRW	CR/UTRW
RI dependents	CHKP (NO)	CHKP (NO)	CHKP (NO)	CHKP (NO)

Legend:
- CHKP (NO): Concurrently running applications will not see CHECK PENDING after commit
- CR: Claim the read claim class
- CW: Claim the write claim class
- DA: Drain all claim classes, no concurrent SQL access
- DR: Drain the repeatable read class, no concurrent access for SQL repeatable readers
- DW: Drain the write claim class, concurrent access for SQL readers
- UTUT: Utility restrictive state, exclusive control
- UTRO: Utility restrictive state, read only access allowed
- UTRW: Utility restrictive state, read/write access allowed
- Blank: Object is not affected by this utility

Compatibility

The following utilities are compatible with LOAD and can run concurrently on the same target object. The target object can be a table space, an index space, or a partition of a table space or index space.

Table 37. Load compatibility

Action	LOAD SHRLEVEL NONE	LOAD SHRLEVEL CHANGE
CHECK DATA DELETE YES	No	No
CHECK DATA DELETE NO	No	No
CHECK INDEX	No	No
CHECK LOB	No	No
COPY INDEXSPACE SHRLEVEL REFERENCE	No	No
COPY INDEXSPACE SHRLEVEL CHANGE	No	Yes
COPY TABLESPACE SHRLEVEL REFERENCE	No	No
COPY TABLESPACE SHRLEVEL CHANGE	No	Yes
COPYTOCOPY	No	Yes

Table 37. Load compatibility (continued)

Action	LOAD SHRLEVEL NONE	LOAD SHRLEVEL CHANGE
DIAGNOSE	Yes	Yes
LOAD SHRLEVEL NONE	No	No
LOAD SHRLEVEL CHANGE	No	Yes
MERGECOPY	No	Yes
MODIFY RECOVERY	No	Yes
MODIFY STATISTICS	No	Yes
QUIESCE	No	No
REBUILD INDEX	No	No
RECOVER (no options)	No	No
RECOVER TOCOPY or TORBA	No	No
RECOVER ERROR RANGE	No	No
REORG INDEX	No	No
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No	No
REORG TABLESPACE UNLOAD ONLY or EXTERNAL	No	No
REPAIR DUMP or VERIFY	No	No
REPAIR LOCATE KEY or RID DELETE or REPLACE	No	No
REPAIR LOCATE TABLESPACE PAGE REPLACE	No	No
REPORT	Yes	No
RUNSTATS TABLESPACE SHRLEVEL REFERENCE	No	No
RUNSTATS TABLESPACE SHRLEVEL CHANGE	No	Yes
RUNSTATS INDEX SHRLEVEL REFERENCE	No	No
RUNSTATS INDEX SHRLEVEL CHANGE	No	Yes
STOSPACE	Yes	Yes
UNLOAD	No	Yes

SQL operations and other online utilities on the same target partition are incompatible.

After running LOAD

The following tasks are described here:

"Copying the loaded table space or partition" on page 209
"Resetting the COPY-pending status" on page 209
"Resetting the REBUILD-pending status" on page 209
"Resetting the CHECK-pending status" on page 210

"Recovering a failed LOAD job" on page 212
"Reorganizing an auxiliary index after LOAD" on page 212

Copying the loaded table space or partition

If you have used LOG YES, consider taking a full image copy of the loaded table space or partition to reduce the processing time of subsequent recovery operations. If you also specified RESUME NO or REPLACE, indicating that this is the first load into the table space, we recommend that you take two or more full image copies to enable recovery. Alternatively, we recommend that you take primary and backup inline copies when you do a LOAD REPLACE; full table space or partition image copies taken after the LOAD completes are not necessary. However, you might need to take images copies of indexes.

Use either the STATISTICS option to collect inline statistics, or the RUNSTATS utility so that the DB2 catalog statistics take into account the newly loaded data, and DB2 can select SQL paths with accurate information. Following this, rebind any application plans that depend on the loaded tables to update the path selection of any embedded SQL statements.

Resetting the COPY-pending status

If you load with LOG NO and do not take an inline copy, LOAD places a table space in the COPY-pending status. Immediately after that operation, DB2 cannot recover the table space (though you can, by loading it again). Prepare for recovery, and turn off the restriction, by making a full image copy using SHRLEVEL REFERENCE. (If you end the copy job before it is finished, the table space is still in COPY-pending status.)

You can also remove the restriction by one of these operations:
- LOAD REPLACE LOG YES
- LOAD REPLACE LOG NO with an inline copy
- REORG LOG YES
- REORG LOG NO with an inline copy
- REPAIR SET with NOCOPYPEND

If you use LOG YES and do not make an image copy of the table space, subsequent recovery operations are possible but will take longer than if you had made an image copy.

A table space in COPY-pending status can be read without restriction; however, it cannot be updated.

Resetting the REBUILD-pending status

LOAD places all the index spaces for a table space in the REBUILD pending status if you end the job (using -TERM UTILITY) before it completes the INDEXVAL phase. It places the table space itself in RECOVER-pending if you end the job before it completes the RELOAD phase.

Resetting the RECOVER-pending status depends on when the utility terminated:

- If the data is intact (running the -DISPLAY DATABASE command shows indexes are in REBUILD-pending status but the table space is not in RECOVER pending status), you can recover the indexes using RECOVER INDEX, if you have a full image copy of the affected indexes. If you do not have an image copy available, you must rebuild the entire index using the REBUILD INDEX utility. However, for

partitioning indexes and for nonpartitioning indexes in REBUILD-pending (RBDP), you can use the PART option of REBUILD INDEX to rebuild separate partitions of the index.

- If the data is not intact (running the -DISPLAY DATABASE command shows the table space is in RECOVER-pending status), you can either load the table again or recover it to a prior point of consistency. The recovery puts the table space into COPY-pending status, and places all indexes in REBUILD-pending status.

Resetting the CHECK-pending status

LOAD places a table space in the CHECK-pending status if its referential integrity is in doubt or its check constraints are violated. The intent of the restriction is to encourage the use of the CHECK DATA utility. That utility locates invalid data and, optionally, removes it. If it removes the invalid data, the data remaining satisfies all check and referential constraints and the CHECK-pending restriction is lifted.

Though CHECK DATA is usually preferred, the CHECK-pending status can also be reset by any of the following operations:

- Dropping tables that contain invalid rows
- Replacing the data in the table space, using LOAD REPLACE and enforcing check and referential constraints
- Recovering all members of the table space set to a prior quiesce point
- REPAIR SET with NOCHECKPEND

In the next sections, we illustrate the use of CHECK DATA after two likely LOAD jobs.

Running CHECK DATA after LOAD REPLACE: Suppose you choose to replace the contents of the project table using LOAD REPLACE. While doing that, you let LOAD enforce its referential and table check constraints, so that the project table contains only valid records at the end of the job; it is *not* in the CHECK-pending status. However, its dependent, the project activity table, *is* placed in CHECK-pending status— some of its rows might have project numbers that are no longer present in the project table. (If the project table had any other dependents, they also would be in CHECK-pending status.)

You want to run CHECK DATA against the table space containing the project activity table to reset the status. First, give particular care to the options described below. Then, when you run the utility, make sure that all table spaces are available that contain either parent tables or dependent tables of any table in the table spaces being checked.

DELETE YES

This option deletes invalid records and resets the status, but it is *not* the default. Use DELETE NO, the default, to find out quickly how large your problem is; you can choose to correct it by reloading, rather than correcting the current situation.

Exception tables

With DELETE YES, you do not use a discard data set to receive copies of the invalid records; instead, you use another DB2 table called an *exception table.* At this point, we assume that you already have an exception table available for every table subject to referential or table check constraints. (For instructions on creating them, see page 58.)

If you use DELETE YES, you must name an exception table for every descendent of every table in every table space being checked. Deletes caused by CHECK DATA are not subject to any of the SQL delete rules; they cascade without restraint to the farthest descendent.

If table Y is the exception table for table X, name it with this clause in the CHECK DATA statement:

```
FOR EXCEPTION IN X USE Y
```

Error and sort data sets

The options ERRDDN, WORKDDN, SORTDEVT, and SORTNUM function in CHECK DATA just as they do in LOAD. That is, you need an error data set, and you can name work data sets for Sort/Merge or let DB2 allocate them dynamically.

The following example runs CHECK DATA against the table space containing the project activity table. It assumes the existence of exception tables named DSN8710.EPROJACT and DSN8710.EEPA.

```
CHECK DATA TABLESPACE DSN8D71A.PROJACT
   DELETE YES
   FOR EXCEPTION IN DSN8710.PROJACT USE DSN8710.EPROJACT
                 IN DSN8710.EMPPROJACT USE DSN8710.EEPA
   SORTDEVT SYSDA
   SORTNUM 4
```

If the statement does not name error or work data sets, the JCL for the job must contain DD statements like these:

```
//SYSERR   DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTWK01 DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTWK02 DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTWK03 DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTWK04 DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//UTPRINT  DD  SYSOUT=A
```

Running CHECK DATA after LOAD RESUME: Suppose now that you want to add records to both the project and project activity tables, using LOAD RESUME. Furthermore, you want to run both jobs at the same time, which you can do because the tables belong to separate table spaces. The only new consideration is that you must load the project activity table using ENFORCE NO, because you cannot assume that the parent project table is already fully loaded.

When the two jobs are complete, what table spaces are in CHECK-pending status?

- If you enforced constraints when loading the project table, it is *not* in CHECK-pending status.
- Because you did not enforce constraints on the project activity table, it *is* in CHECK-pending status.
- Because you used LOAD RESUME (not LOAD REPLACE) when loading the project activity table, its dependents (the employee to project activity table) are *not* in CHECK-pending status. That is, the operation might not delete any parent rows from the project table, and so might not violate the referential integrity of its dependent. But if you delete records from PROJACT when checking, you still need an exception table for EMPPROJACT.

Hence you want to check the data in the project activity table.

SCOPE PENDING

DB2 records the identifier of the first record of the table that might violate referential or table check constraints. For partitioned table spaces, that identifier is in SYSIBM.SYSTABLEPART; for nonpartitioned table spaces, that identifier is in SYSIBM.SYSTABLES. The SCOPE PENDING option speeds the checking by confining it to just the records that might be in error.

The following example runs CHECK DATA against the table space containing the project activity table after LOAD RESUME:

```
CHECK DATA TABLESPACE DSN8D71A.PROJACT
   SCOPE PENDING
   DELETE YES
   FOR EXCEPTION IN DSN8710.PROJACT USE DSN8710.EPROJACT
                 IN DSN8710.EMPPROJACT USE DSN8710.EEPA
   SORTDEVT SYSDA
   SORTNUM 4
```

As before, the JCL data set for the job needs DD statements to define data sets for the error and sort data sets.

Collecting inline statistics while loading a table

If you do not specify LOAD RESUME, use the STATISTICS keyword to gather inline statistics about space use and row clustering to update the DB2 catalog. The data is used to select access paths when executing SQL statements. This procedure eliminates the need to run RUNSTATS after loading a table space. However, it you perform a LOAD PART operation, you should run RUNSTATS INDEX on the nonpartitioning indexes to update that data.

Running CHECK INDEX after loading a table having indexes

The CHECK INDEX utility tests whether an index is consistent with the data it indexes and issues error messages if it finds an inconsistency. If you have any reason to doubt the accuracy of an index (for example, if the result of an SQL SELECT COUNT statement is inconsistent with the output of RUNSTATS), run CHECK INDEX. You might also want to run CHECK INDEX after any LOAD operation that shows some abnormal condition in its execution, or even run it periodically to verify the accuracy of important indexes.

To rebuild an index that is inconsistent with its data, use the REBUILD INDEX utility.

Recovering a failed LOAD job

To facilitate recovery in case of failure, the SYSCOPY record is inserted at the beginning of the RELOAD phase if LOG YES was specified in the LOAD control statement. As a result, you can recover the data to a point in time before the LOAD by using RECOVER TORBA.

Reorganizing an auxiliary index after LOAD

Indexes on the auxiliary tables are not built during the BUILD phase. Instead, LOB values are inserted (not loaded) into auxiliary tables during the RELOAD phase as each row is loaded into the base table, and each index on the auxiliary table is updated as part of the INSERT operation. Because the LOAD utility inserts keys into an auxiliary index, free space within the index might be consumed and index page splits might occur. Consider reorganizing an index on the auxiliary table after LOAD completes to introduce free space into the index for future INSERTs and LOADs.

Sample control statements

Example 1: LOAD JCL with RESUME YES and ENFORCE NO. This example shows the JCL for loading a table with the RESUME YES and ENFORCE NO options. This job will place the table in the CHECK-pending status.

```
//STEP1    EXEC DSNUPROC,UID='IUIQU2UB.LOAD',
//         UTPROC='',
//         SYSTEM='V71A'
//SYSRECAC   DD DSN=IUIQU2UB.LOAD.DATA,DISP=SHR,VOL=SER=SCR03,
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD DSN=IUIQU2UB.LOAD.STEP1.SYSUT1,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD DSN=IUIQU2UB.LOAD.STEP1.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
LOAD DATA INDDN(SYSRECAC) RESUME YES
     INTO TABLE DSN8710.ACT
          (ACTNO    POSITION( 1) INTEGER EXTERNAL(3),
           ACTKWD   POSITION( 5) CHAR(6),
           ACTDESC  POSITION(13) VARCHAR)
     ENFORCE NO
//*
```

Example 2: Control statement with RESUME YES option. Figure 13 shows a LOAD utility statement. It loads the records from the data set named by the SYSREC DD statement for the utility job into the department table. The RESUME YES clause specifies that the table space need not be empty; new records are added at the end.

```
LOAD DATA
  RESUME YES
  INTO TABLE DSN8710.DEPT
  ( DEPTNO    POSITION (1:3)     CHAR(3),
    DEPTNAME  POSITION (4:39)    CHAR(36),
    MGRNO     POSITION (40:45)   CHAR(6),
    ADMRDEPT  POSITION (46:48)   CHAR(3),
    LOCATION  POSITION (49:64)   CHAR(16) )
```

Figure 13. Example of a LOAD utility statement

This example uses the POSITION clause to specify where a field is in the input record. With the statement above, LOAD accepts the input shown in Figure 14 on page 214 and interprets it as follows:

- The first three bytes of each record are loaded into the DEPTNO column of the table.
- The next 36 bytes are loaded into the DEPTNAME column, including trailing blanks.

  If we had chosen to define this input column as VARCHAR(36), the input data would have had to contain a 2-byte binary length field preceding the data.
- The next three fields are loaded into columns defined as CHAR(6), CHAR(3), and CHAR(16).

```
A00SPIFFY COMPUTER SERVICE DIV.        000010A00USIBMSTODB21
B01PLANNING                            000020A00USIBMSTODB21
C01INFORMATION CENTER                  000030A00USIBMSTODB21
D01DEVELOPMENT CENTER                        A00USIBMSTODB21
```

Figure 14. Records in an input data set for LOAD

Table 38 shows how the same records appear if you then execute the statement SELECT * FROM DSN8710.DEPT under SPUFI.

Table 38. Data loaded to a table

DEPTNO	DEPTNAME	MGRNO	ADMRDEPT	LOCATION
A00	SPIFFY COMPUTER SERVICE DIV.	000010	A00	USIBMSTODB21
B01	PLANNING	000020	A00	USIBMSTODB21
C01	INFORMATION CENTER	000030	A00	USIBMSTODB21
D01	DEVELOPMENT CENTER	______	A00	USIBMSTODB21

Example 3: Load data into a table. Load data from the data set specified by the EMPLDS DD statement into the EMP table.

```
LOAD DATA INDDN EMPLDS
  INTO TABLE DSN8710.EMP
```

Example 4: Load data into two tables. Load data from the data set specified by the EMPLDS DD statement into the DSN8710.EMP and SMITH.EMPEMPL tables.

```
LOAD DATA INDDN EMPLDS
  INTO TABLE DSN8710.EMP
  INTO TABLE SMITH.EMPEMPL
```

Example 5: Load selected records into a table. Load data from the data set specified by the EMPLDS DD statement into the EMP table. Load only from source input records that begin with LKA.

```
LOAD DATA INDDN EMPLDS
  INTO TABLE DSN8710.EMP
  WHEN (1:3)='LKA'
```

Example 6: Load selected records into a non-empty table space. The data from the sequential data set identified by the SYSREC DD statement is selectively loaded into the DSN8710.DEPT table whenever positions 1 through 3 contain the value LKA. The table space need not be empty for loading to proceed.

For each source record that has LKA in its first three positions:

- The characters in positions 7 through 9 are loaded into the DEPTNO column
- The characters in positions 10 through 35 are loaded into the DEPTNAME VARCHAR column
- The characters in positions 36 through 41 are loaded into the MGRNO column
- Characters in positions 42 through 44 are loaded into the ADMRDEPT column.

```
LOAD DATA
  RESUME YES
  INTO TABLE DSN8710.DEPT WHEN (1:3)='LKA'
  (DEPTNO POSITION (7:9) CHAR,
   DEPTNAME POSITION (10:35) CHAR,
   MGRNO POSITION (36:41) CHAR,
   ADMRDEPT POSITION (42:44) CHAR)
```

Example 7: Load selected records into an empty table space. Data from the sequential data set identified by the SYSRECPJ DD statement is selectively loaded into the DSN8710.PROJ table. The table space containing the DSN8710.PROJ table is currently empty, because the RESUME YES option was not specified.

For each source input record, data is loaded into the specified columns (that is, PROJNO, PROJNAME, DEPTNO ..., and so on) to form a table row. Any other columns in a DSN8710.PROJ row are set to NULL.

Starting positions of the fields in the sequential data set are defined by the field specification POSITION options. The ending position of the fields in the sequential data set are implicitly defined either by the length specification of the data type options (CHAR length) or by the length specification of the external numeric data type (LENGTH).

The numeric data represented in SQL constant format (EXTERNAL format) is converted to the correct internal format by the LOAD process and placed in the indicated column names. The two dates are assumed to be represented by eight digits and two separator characters, as in the USA format (for example, 11/15/1987). The length of the date fields is given as 10 explicitly, though in many cases it defaults to the same value.

```
LOAD DATA INDDN(SYSRECPJ)
  INTO TABLE DSN8710.PROJ
  (PROJNO    POSITION   (1) CHAR(6),
  PROJNAME   POSITION   (8) CHAR(22),
  DEPTNO     POSITION  (31) CHAR(3),
  RESPEMP    POSITION  (35) CHAR(6),
  PRSTAFF    POSITION  (42) DECIMAL EXTERNAL(5),
  PRSTDATE   POSITION  (48) DATE EXTERNAL(10),
  PRENDATE   POSITION  (59) DATE EXTERNAL(10),
  MAJPROJ    POSITION  (70) CHAR(6))
```

Example 8: Load data selectively using the CONTINUEIF option. Data from the sequential data set specified by the SYSRECOV DD statement is assembled and selectively loaded into the DSN8710.TOPTVAL table. The table space that contains DSN8710.TOPTVAL is currently empty because the RESUME YES option is not specified.

Fields destined for columns in the same table row can span more than one source record. Source records having fields containing columns that belong to the same row as the next source record all have an X in column 72 (that is, `CONTINUEIF(72:72)='X'`).

For each assembled source record, fields are loaded into the DSN8710.TOPTVAL table columns (that is, MAJSYS, ACTION, OBJECT ..., DSPINDEX) to form a table row. Any columns not mentioned are set to NULL.

The starting positions of the fields in the assembled source record input are given in the POSITION option. Starting positions are numbered from the first column of the internally assembled input record, not from the start of the source records in the sequential data set. The ending positions are defined by the character string lengths given with the input data type.

No conversions are required to load the source character strings into their designated columns, which are also defined to be fixed character strings. However,

because columns INFOTXT, HELPTXT, and PFKTXT are defined as 79 characters in length and the strings being loaded are 71 characters in length, those strings are padded with blanks as they are loaded.

```
LOAD DATA INDDN(SYSRECOV) CONTINUEIF(72:72)='X'
  INTO TABLE DSN8710.TOPTVAL
  (MAJSYS   POSITION   (2) CHAR(1),
   ACTION   POSITION   (4) CHAR(1),
   OBJECT   POSITION   (6) CHAR(2),
   SRCHCRIT POSITION   (9) CHAR(2),
   SCRTYPE  POSITION  (12) CHAR(1),
   HEADTXT  POSITION  (80) CHAR(50),
   SELTXT   POSITION (159) CHAR(50),
   INFOTXT  POSITION (238) CHAR(71),
   HELPTXT  POSITION (317) CHAR(71),
   PFKTXT   POSITION (396) CHAR(71),
   DSPINDEX POSITION (475) CHAR(2))
```

Example 9: Load data with referential constraints. Data from the sequential data set identified by the SYSREC DD statement is loaded into the DSN8710.PROJ. table. Referential constraints are enforced on data added. Output consists of a summary report of violations of referential constraints, and all records causing these violations are placed in the SYSDISC discard data set.

```
LOAD DATA INDDN(SYSREC) CONTINUEIF(72:72)='X'
 RESUME YES
 ENFORCE CONSTRAINTS
 INTO TABLE DSN8710.PROJ
  (PROJNO POSITION (1) CHAR (6),
   PROJNAME POSITION (8) VARCHAR,
   DEPTNO POSITION (33) CHAR (3),
   RESPEMP POSITION (37) CHAR (6),
   PRSTAFF POSITION (44) DECIMAL EXTERNAL (5),
   PRSTDATE POSITION (50) DATE EXTERNAL,
   PRENDATE POSITION (61) DATE EXTERNAL,
   MAJPROJ POSITION (80) CHAR (6) NULLIF(MAJPROJ='      '))
```

Example 10: Load data using SORTKEYS. Use the SORTKEYS keyword to improve performance of the index key sort as shown in the following example. Assume there are 22,000 rows to load into the DSN8710.DEPT table. This table has 3 indexes.

The following job specifies an estimate of 66,000 keys to sort with the SORTKEYS keyword, using the calculation described in “Improving performance with SORTKEYS” on page 196:

```
(3 + 0) * 22,000 = 66,000
```

This example specifies dynamic allocation of the required data sets by DFSORT, using the SORTDEVT and SORTNUM keywords. If sufficient virtual storage resources are available, one utility subtask pair will be started to build each index. This example does not require UTPRIN*nn* DD statements, because it uses DSNUPROC to invoke utility processing, which includes a DD statement that allocates UTPRINT to SYSOUT.

LOAD statement:

```
//SAMPJOB  JOB  ...
//STEP1    EXEC DSNUPROC,UID='SAMPJOB.LOAD',UTPROC='',SYSTEM='DSN'
//SORTOUT  DD DSN=SAMPJOB.LOAD.STEP1.SORTOUT,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SYSUT1   DD DSN=SAMPJOB.LOAD.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SYSERR   DD DSN=SAMPJOB.LOAD.STEP1.SYSERR,DISP=(MOD,DELETE,CATLG),
```

```
//         UNIT=SYSDA,SPACE=(2000,(20,20),,,ROUND)
//         DCB=(RECFM=FB,LRECL=80,BLKSIZE=2400)
//SYSMAP   DD DSN=SAMPJOB.LOAD.STEP1.SYSMAP,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(2000,(20,20),,,ROUND),
//         DCB=(RECFM=FB,LRECL=80,BLKSIZE=2400)
//SYSREC   DSN=SAMPJOB.TEMP.DATA,DISP=SHR,UNIT=SYSDA
//SYSIN    DD *
LOAD DATA REPLACE INDDN SYSREC CONTINUEIF(79:80)='++'
 SORTKEYS 66000 SORTDEVT SYSDA SORTNUM 3
 INTO TABLE DSN8710.DEPT
/*
```

Example 11: LOAD with inline copy. Use the REPLACE option with COPYDDN and RECOVERYDDN to create copies during LOAD.

```
LOAD DATA REPLACE INDDN INPUT
 SORTKEYS 66000
 COPYDDN SYSCOPY RECOVERYDDN REMCOPY
 CONTINUEIF(79:80)='++'
 INTO TABLE DSN8710.DEPT
```

Example 12: Load ASCII input data. Use the ASCII option to load ASCII input data into a table named MYASCIIT that was created with the CCSID ASCII clause.

```
LOAD REPLACE LOG NO ASCII INTO TABLE MYASCIIT
 (NAME POSITION(1) CHAR(40),
  ADDRESS POSITON(41) CHAR(40),
  ZIP POSITION(81) DECIMAL EXTERNAL(9),
  DEPARTMENT POSITION(90) CHAR(3),
  TITLE POSITION(93) GRAPHIC(20))
```

The CCSID keyword is not specified in this example; therefore, the CCSIDs of the ASCII input data are assumed to be the ASCII CCSIDs specified at installation. Conversions are done only if the CCSIDs of the target table differ from the ASCII CCSIDs specified at installation.

Example 13: Load data using statistics collection. Use the STATISTICS keyword to gather catalog statistics for the table space. This eliminates the need to run the RUNSTATS utility after completing the load operation. Specify REUSE so that all partitions are logically reset rather than deleted and redefined.

```
LOAD DATA
     INDDN SYSREC
     REPLACE
     STATISTICS TABLE(ALL)
     INDEX(ALL)
     REPORT YES UPDATE ALL
     REUSE
     CONTINUEIF(79:80)='++'
     INTO TABLE
     DSN8710.DEPT
```

Example 14: Load data for a partitioned table space using statistics collection. Load data for a specified partition, using the STATISTICS keyword to gather catalog statistics of the partitioned table space.

```
LOAD STATISTICS
  INTO TABLE DSN8710.DEPT PART 1 REPLACE
```

Example 15: Load data from individual input data sets into the partitions using Partition Parallelism, with SHRLEVEL CHANGE access. Use the TEMPLATE control statement to define the data set naming convention for dynamic allocation.

```
//LOAD3    JOB  ...
//STEP1    EXEC DSNUPROC,UID='JULXU333.LOAD3',
//         UTPROC='',
//         SYSTEM='V71A'
//UTPRINT  DD  SYSOUT=*
//PART1    DD DSN=JULXU333.TBL3.PART1,
//         DISP=(OLD,KEEP,CATLG),UNIT=SYSDA
//PART2    DD DSN=JULXU333.TBL3.PART2,
//         DISP=(OLD,KEEP,CATLG),UNIT=SYSDA
//PART3    DD DSN=JULXU333.TBL3.PART3,
//         DISP=(OLD,KEEP,CATLG),UNIT=SYSDA
//PART4    DD DSN=JULXU333.TBL3.PART4,
//         DISP=(OLD,KEEP,CATLG),UNIT=SYSDA
//PART5    DD DSN=JULXU333.TBL3.PART5,
//         DISP=(OLD,KEEP,CATLG),UNIT=SYSDA
//DISC1    DD DSN=JULXU333.LOAD3.TBL3.DISC1,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//DISC2    DD DSN=JULXU333.LOAD3.TBL3.DISC2,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//DISC3    DD DSN=JULXU333.LOAD3.TBL3.DISC3,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//DISC4    DD DSN=JULXU333.LOAD3.TBL3.DISC4,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//DISC5    DD DSN=JULXU333.LOAD3.TBL3.DISC5,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD DSN=JULXU333.LOAD3.STEP1.SYSUT1,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD DSN=JULXU333.LOAD3.STEP1.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
  TEMPLATE ERR3
         DSN &UT..&JO..&ST..ERR3&MO.&DAY.
         UNIT SYSDA  DISP(NEW,CATLG,CATLG)
  LOAD DATA
    LOG YES  RESUME YES  SHRLEVEL CHANGE
    ERRDDN ERR3
    INTO TABLE DBA01.TBLX3303
      PART 1
      INDDN PART1
      DISCARDDN DISC1
        (EMPNO      POSITION(1)    CHAR(6),
         LASTNAME   POSIITON(8)    VARCHAR(15),
         SALARY     POSITION(25)   DECIMAL(9,2))
    .
    .
    .
    INTO TABLE DBA01.TBLX3303
      PART 5
      INDDN PART5
      DISCARDDN DISC5
        (EMPNO      POSITION(1)    CHAR(6),
         LASTNAME   POSIITON(8)    VARCHAR(15),
         SALARY     POSITION(25)   DECIMAL(9,2))
/*
```

Example 16: Load data using a declared cursor. Load MYEMP table with the results of the SELECT statement declared for cursor C1.

```
EXEC SQL
  DECLARE C1 CURSOR FOR SELECT * FROM DSN8710.EMP
ENDEXEC

LOAD DATA
INCURSOR(C1)
REPLACE
INTO TABLE MYEMP
STATISTICS
```

Example 17: Load data partitions in parallel using a declared cursor. Load the partitions of MYEMP table in parallel with the results of the SELECT statement declared for each partition.

```
EXEC SQL
  DECLARE C1 CURSOR FOR SELECT * FROM CHICAGO.DSN8710.EMP
  WHERE EMPNO <= '099999'
ENDEXEC
EXEC SQL
  DECLARE C2 CURSOR FOR SELECT * FROM CHICAGO.DSN8710.EMP
  WHERE EMPNO > '099999' AND EMPNO <= '199999'
ENDEXEC
EXEC SQL
  DECLARE C3 CURSOR FOR SELECT * FROM CHICAGO.DSN8710.EMP
  WHERE EMPNO > '199999' AND EMPNO <= '299999'
ENDEXEC
EXEC SQL
  DECLARE C4 CURSOR FOR SELECT * FROM CHICAGO.DSN8710.EMP
  WHERE EMPNO > '299999' AND EMPNO <= '999999'
ENDEXEC

LOAD DATA
  INTO TABLE MYEMPP PART 1 REPLACE INCURSOR(C1)
  INTO TABLE MYEMPP PART 2 REPLACE INCURSOR(C2)
  INTO TABLE MYEMPP PART 3 REPLACE INCURSOR(C3)
  INTO TABLE MYEMPP PART 4 REPLACE INCURSOR(C4)
```

Chapter 15. MERGECOPY

The MERGECOPY online utility merges image copies produced by the COPY utility or inline copies produced by the LOAD or REORG utilities. It can merge several incremental copies of a table space to make one incremental copy. It can also merge incremental copies with a full image copy to make a new full image copy.

MERGECOPY operates on the image copy data sets of a table space, and not on the table space itself.

For a diagram of MERGECOPY syntax and a description of available options, see "Syntax and options of the control statement" on page 222. For detailed guidance on running this utility, see "Instructions for running MERGECOPY" on page 224.

Output: Output from the MERGECOPY utility consists of one of the following types of copies:
- A new single incremental image copy
- A new full image copy

You can create the new image copy for the local or recovery site.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- IMAGCOPY privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority

An ID with installation SYSOPR authority can also execute MERGECOPY, but only on a table space in the DSNDB01 or DSNDB06 database.

Restrictions on running MERGECOPY:
- MERGECOPY cannot merge image copies into a single incremental image copy for the other site, that is,
  - At local sites, you cannot use RECOVERYDDN with NEWCOPY NO.
  - At recovery sites, you cannot use COPYDDN with NEWCOPY NO.
- When none of the keywords NEWCOPY, COPYDDN, or RECOVERYDDN is specified, the default, NEWCOPY NO COPYDDN(SYSCOPY), is valid for the local site only.

Execution phases of MERGECOPY: One of the following phases can be identified if the job terminates.

The phases for MERGECOPY are:

Phase	Description
UTILINIT	Initialization
MERGECOP	Merge incremental copies
UTILTERM	Cleanup.

© Copyright IBM Corp. 1983, 2001

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, you can use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

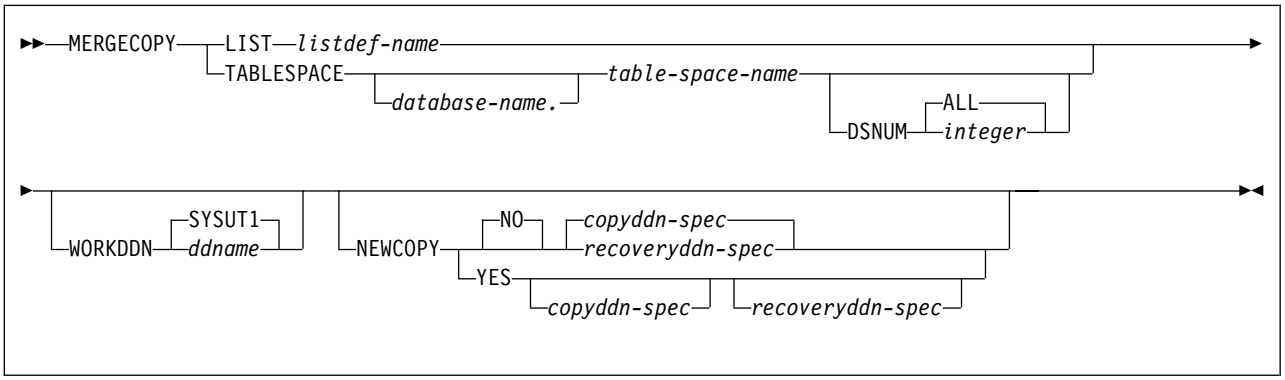

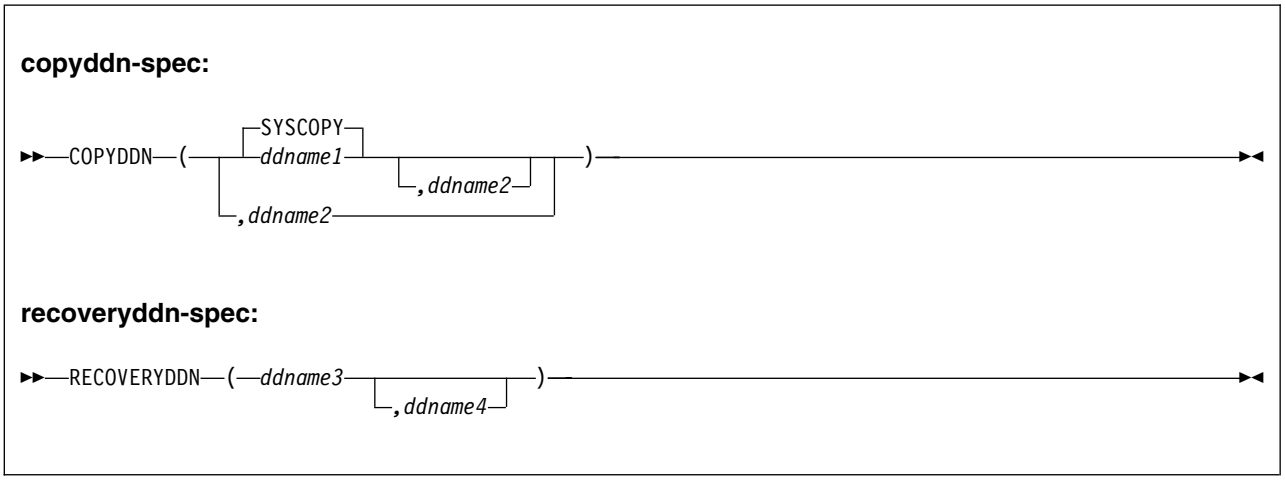

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

LIST *listdef-name*

Specifies the name of a previously defined LISTDEF list name. LIST specifies one LIST keyword for each COPY control statement. Do not specify LIST with either the INDEX keywords or the TABLESPACE, utility list manager will invoke COPY once for the entire list. For more information about LISTDEF specifications, see "Generating object lists using LISTDEF" on page 144.

TABLESPACE *database-name.table-space-name*
Specifies the table space (and, optionally, the database it belongs to) that is to be copied.

database-name
The name of the database the table space belongs to. The **default** is **DSNDB04**.

table-space-name
The name of the table space whose incremental image copies are to be merged.

You cannot run the MERGECOPY utility on the DSNDB01.DBD01, DSNDB01.SYSUTILX, or DSNDB06.SYSCOPY table spaces because you cannot make incremental copies of those table spaces. Because MERGECOPY does not directly access the table space whose copies it is merging, it does not interfere with concurrent access to that table space.

The following are optional.

WORKDDN *ddname*
Specifies a DD statement for a temporary data set or template, to be used for intermediate merged output.

ddname is the DD name. The **default** is **SYSUT1**.

Use the WORKDDN option if you are not able to allocate enough data sets to execute MERGECOPY; in that case, a temporary data set is used to hold intermediate output. If you omit the WORKDDN option, it is possible that only some of the image copy data sets will be merged. When MERGECOPY has ended, a message is issued that tells the number of data sets that exist and the number of data sets that have been merged. To continue the merge, repeat MERGECOPY with a new output data set.

DSNUM
Identifies a partition or data set, within the table space, that is to be merged; or it merges the entire table space.

ALL Merges the entire table space. The **default** is **ALL**.

integer
Is the number of a partition or data set to be merged. The maximum is 254.

For a partitioned table space, the integer is its partition number.

For a nonpartitioned table space, find the integer at the end of the data set name as cataloged in the VSAM catalog. The data set name has this format:

catname.DSNDB*x*.*dbname*.*tsname*.*y*0001.A*nnn*

where *y* is either I or J and *nnn* is the data set integer.

This variable may not be specified with LIST. Use PARTLEVEL on the LISTDEF instead. If image copies were taken by data set (rather than by table space), then MERGECOPY must use the copies by data set.

NEWCOPY
Tells whether incremental image copies are to be merged with the full image copy or not.

NO
Merges incremental image copies into a single incremental image copy, but does not merge them with the full image copy. The **default** is **NO**.

YES
Merges all incremental image copies with the full image copy to form a new full image copy.

COPYDDN (*ddname1,ddname2***)**
Specifies the DD statements for the output image copy data sets at the local site. *ddname1* is the primary output image copy data set. *ddname2* is the backup output image copy data set.

The **default** is **COPYDDN**(**SYSCOPY**), where SYSCOPY identifies the primary data set.

The COPYDDN keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

RECOVERYDDN (*ddname3,ddname4***)**
Specifies the DD statements for the output image copy data sets at the recovery site. You can have a maximum of two output data sets; the outputs are identical. *ddname3* is the primary output image copy data set. *ddname4* is the backup output image copy data set.

There is no default for RECOVERYDDN.

The **RECOVERYDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

Instructions for running MERGECOPY

To run MERGECOPY, you must:

1. Prepare the necessary data sets, as described in "Data sets used by MERGECOPY" on page 225.
2. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for MERGECOPY, see "Sample control statements" on page 229.)
3. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 226. (For a complete description of the syntax and options for MERGECOPY, see "Syntax and options of the control statement" on page 222.)
4. Check the compatibility table in "Concurrency and compatibility" on page 228 if you want to run other jobs concurrently on the same target objects.
5. Plan for restart if the MERGECOPY job doesn't complete, as described in "Terminating or restarting MERGECOPY" on page 228.
6. Run MERGECOPY.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Data sets used by MERGECOPY

Table 39 describes the data sets used by MERGECOPY. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 39. Data sets used by MERGECOPY

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
Image copy data set	The image copy data set defines the resulting image copy. Its DD names is specified with the COPYDDN parameter of the MERGECOPY statement. The default DD name is SYSCOPY.	Yes
Work data set	This is a temporary data set that is used for intermediate merged output. Its DD name is specified through the WORKDDN parameter of the MERGECOPY statement. The default DD name is SYSUT1.	Yes
Input data sets	These are image copy data sets which may be pre-allocated by the user. The DD names are defined by the user.	No

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table space
Object to be copied. It is named in the MERGECOPY control statement and is accessed through the DB2 catalog.

Data sets: The input data sets for the merge operation are dynamically allocated. To merge incremental copies, a work data set (WORKDDN) and up to two new copy data sets (COPYDDN) are allocated in the JCL for the utility job. These can be allocated to tape or DASD. If these allocations are made to tape, an additional tape drive is required for each of those data sets.

The COPYDDN option of MERGECOPY allows you to specify the ddnames for the output data sets. The option has the format `COPYDDN (ddname1,ddname2)`, where *ddname1* is the ddname for the primary output data set in the system currently running DB2 and *ddname2* is the ddname for the backup output data set in the system currently running DB2. The default for *ddname1* is SYSCOPY.

The RECOVERYDDN option of MERGECOPY allows you to specify the output image copy data sets at the recovery site. The option has the format `RECOVERYDDN (ddname3, ddname4)`, where *ddname3* is the ddname for the primary output image copy data set at the recovery site and *ddname4* is the ddname for the backup output data set at the recovery site.

Defining the work data set: The work data set should be at least equal in size to the largest input image copy data set being merged. Use the same DCB attributes as used for the image copy data sets.

Creating the control statement

See “Syntax and options of the control statement” on page 222 for MERGECOPY syntax and option descriptions. See “Sample control statements” on page 229 for examples of MERGECOPY usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

“Specifying full or incremental image copy”
“Merging inline copies” on page 227
“Using MERGECOPY with individual data sets” on page 227
“Deciding between MERGECOPY or COPY” on page 227
“Avoiding MERGECOPY LOG RBA inconsistencies” on page 227
“Creating image copies in a JES3 environment” on page 228
“Running MERGECOPY on the directory” on page 228

Specifying full or incremental image copy

The NEWCOPY parameter decides if the new copy created by MERGECOPY is an incremental image copy or a full image copy. In general, it is recommended to create a new full image copy. The reasons for this recommendation are:

- A new full image copy creates a new recoverable point.
- The additional time it takes to create a new full image copy does not have any adverse affect on the access to the table space. The only concurrency implication is the access to SYSIBM.SYSCOPY.
- The range of log records to be applied by RECOVER is the same for both the new full image copy and the merged incremental image copy.
- Assuming the copies are on tape, only one tape drive will be required for image copies during a RECOVER.

If NEWCOPY is YES, the utility inserts an entry for the new full image copy into the SYSIBM.SYSCOPY catalog table.

If NEWCOPY is NO, the utility:

- Replaces the SYSIBM.SYSCOPY records of the incremental image copies that were merged with an entry for the new incremental image copy
- Deletes all SYSIBM.SYSCOPY records of the incremental image copies that have been merged.

In either case, if any of the input data sets might not be allocated, or you did not specify a temporary work data set (WORKDDN), the utility performs a partial merge.

For large table spaces, the use of MERGECOPY to create full image copies should be considered.

Use MERGECOPY NEWCOPY YES immediately after each incremental image copy:

- Dates will become a valid criterion for image copy data set and archive log deletion.
- A minimum number of tape drives will be allocated for MERGECOPY and RECOVER execution.

Merging inline copies

If you merge an inline copy with incremental copies, the result is a full inline copy. The data set is logically equivalent to a full image copy, but the data within the data set differs in some respects. See “Using inline COPY with LOAD” on page 194 for additional information about inline copies.

Using MERGECOPY with individual data sets

MERGECOPY can be used on copies of an entire table space or individual data sets or partitions. However, MERGECOPY can only merge incremental copies of the same type. That is, you cannot merge incremental copies of an entire table space with incremental copies of individual data sets to form new incremental copies. The attempt to mix the two types of incremental copies produces the following messages:

```
DSNU460I DSNUBCLO - IMAGE COPIES INCONSISTENT.
                    MERGECOPY REQUEST REJECTED
DSNU010I DSNUGBAC - UTILITY EXECUTION COMPLETE,
                    HIGHEST RETURN CODE=4
```

With the option NEWCOPY YES, however, you can merge a full image copy of a table space with incremental copies of the table space and of individual data sets to make a new full image copy of the table space.

If the image copy data sets you want to merge reside tape, refer to “Retaining tape mounts” on page 298 for general information about specifying the appropriate parameters on the DD cards.

Deciding between MERGECOPY or COPY

COPY and MERGECOPY can create a full image copy. COPY is required after a LOAD or REORG with LOG NO unless an inline copy is created, but in other cases an incremental image copy followed by MERGECOPY is a valid alternative.

Avoiding MERGECOPY LOG RBA inconsistencies

MERGECOPY does not use information that was logged between the time of the most recent image copy and the time when MERGECOPY was run. Therefore, you cannot safely delete all log records made before running MERGECOPY. (And you do that if you run MODIFY RECOVERY specifying the date when MERGECOPY was run as the value of DATE.)

To delete all log information that is included in a copy made by MERGECOPY:

1. Find the record of that copy in the catalog table SYSIBM.SYSCOPY. You can find it by selecting on database name, table space name, and date (columns DBNAME, TSNAME, and ICDATE).
2. Column START_RBA contains the RBA of the last image copy that MERGECOPY used. Find the record of the image copy that has the same value of START_RBA.
3. In that record, find the date in column ICDATE. You can use MODIFY RECOVERY to delete all copies and log records for the table space made before that date.

RECOVER uses the LOG RBA of image copies to determine the starting point in the log needed for recovery. Normally, there is a direct correspondence between a timestamp and a LOG RBA. Because of this, and because MODIFY uses dates to cleanup recovery history, you may decide to use dates to delete old archive log tapes. This may cause a problem if MERGECOPY is used. MERGECOPY inserts the LOG RBA of the last incremental image copy into the SYSIBM.SYSCOPY row

created for the new image copy. The date recorded in ICDATE column of SYSIBM.SYSCOPY row is the date MERGECOPY was executed.

Creating image copies in a JES3 environment

Ensure that there are sufficient units available to mount the required image copies. In a JES3 environment, if the number of image copies to be restored exceeds the number of available online and offline units, and the MERGECOPY job successfully allocates all available units, the job will then wait for more units to become available.

Running MERGECOPY on the directory

You cannot run the MERGECOPY utility on the DSNDB01.DBD01, DSNDB01.SYSUTILX, or DSNDB06.SYSCOPY table spaces because you cannot make incremental copies of those table spaces.

Terminating or restarting MERGECOPY

For instructions on restarting a utility job, see “Restarting an online utility” on page 46.

You can restart a MERGECOPY utility job at the beginning of any of the phases listed below:

UTILINIT
Initialization and setup

MERGECOPY
Merge

UTILTERM
Cleanup.

Restarting MERGECOPY after an out of space condition: See “Restarting after the output data set is full” on page 47 for guidance in restarting MERGECOPY from the last commit point after receiving an out of space condition.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Table 40 shows the restrictive state the utility sets on the target object.

Table 40. Claim classes of MERGECOPY operations. Use of claims and drains; restrictive states set on the target object.

Target	MERGECOPY
Table space or partition	UTRW

Legend:

UTRW - Utility restrictive state - Read/Write access allowed.

MERGECOPY can run concurrently on the same target object with any utility **except the following**:

- COPY TABLESPACE
- LOAD
- MERGECOPY
- MODIFY
- RECOVER

- REORG TABLESPACE
- UNLOAD (only when from the same image copy data set)

The target object can be a table space or partition.

Sample control statements

Example 1: Creating a merged incremental copy. Create a merged incremental image copy of table space DSN8S71C.

```
//STEP1    EXEC DSNUPROC,UID='IUJMU107.MERGE1',
//          UTPROC='',SYSTEM='V71A'
//COPY1  DD DSN=IUJMU107.MERGE1.STEP1.COPY1,DISP=(MOD,CATLG,CATLG),
//          UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//COPY2  DD DSN=IUJMU107.MERGE1.STEP1.COPY2,DISP=(MOD,CATLG,CATLG),
//          UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD DSN=IUJMU107.MERGE1.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//          UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *

MERGECOPY TABLESPACE DSN8D71P.DSN8S71C
  COPYDDN      (COPY1,COPY2)
  NEWCOPY  NO
```

Example 2: Creating a merged full image copy. Create a merged full image copy of table space DSN8S71C.

```
//STEP1    EXEC DSNUPROC,UID='IUJMU107.MERGE2',
//          UTPROC='',SYSTEM='DSN'
//COPY1  DD DSN=IUJMU107.MERGE2.STEP1.COPY1,DISP=(MOD,CATLG,CATLG),
//          UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//COPY2  DD DSN=IUJMU107.MERGE2.STEP1.COPY2,DISP=(MOD,CATLG,CATLG),
//          UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD DSN=IUJMU107.MERGE2.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//          UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
MERGECOPY TABLESPACE DSN8D71P.DSN8S71C
  COPYDDN      (COPY1,COPY2)
  NEWCOPY  YES
```

Chapter 16. MODIFY RECOVERY

The MODIFY online utility with the RECOVERY option deletes records from the SYSIBM.SYSCOPY catalog table, related log records from the SYSIBM.SYSLGRNX directory table, and entries from the DBD. You can remove records that were written before a specific date or you can remove records of a specific age. You can delete records for an entire table space, partition, or data set.

You should run MODIFY regularly to clear outdated information from SYSIBM.SYSCOPY and SYSIBM.SYSLGRNX. These tables, and particularly SYSIBM.SYSLGRNX, can become very large and take up considerable amounts of space. By deleting outdated information from these tables, you can help improve performance for processes that access data from these tables.

The MODIFY RECOVERY utility automatically removes the SYSIBM.SYSCOPY and SYSIBM.SYSLGRNX recovery records that meet the AGE and DATE criteria for all indexes over the table space that were defined with the COPY YES attribute.

For a diagram of MODIFY RECOVERY syntax and a description of available options, see “Syntax and options of the control statement” on page 232. For detailed guidance on running this utility, see “Instructions for running MODIFY RECOVERY” on page 233.

Output: MODIFY RECOVERY deletes image copy rows from SYSIBM.SYSCOPY and SYSIBM.SYSLGRNX.

For each full and incremental SYSCOPY record deleted from SYSCOPY, the utility returns a message giving the name of the copy data set.

For information on deleting SYSLGRNX rows, see “Deleting SYSLGRNX rows” on page 234.

If MODIFY RECOVERY deletes at least one SYSCOPY record and the target table space or partition is not recoverable, the target object is placed in COPY-pending status.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- IMAGCOPY privilege for the database to run MODIFY RECOVERY
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority.

An ID with installation SYSOPR authority can also execute MODIFY RECOVERY, but only on a table space in the DSNDB01 or DSNDB06 database.

There are no SYSCOPY or SYSLGRNX records for DSNDB06.SYSCOPY, DSNDB01.SYSUTIL, or DSNDB01.DBD01. You can run MODIFY RECOVERY on these table spaces, but you receive message DSNU573I, indicating that no SYSCOPY records could be found. No SYSCOPY or SYSLGRNX records are deleted.

Execution phases of MODIFY RECOVERY: One of the following phases can be identified if the job terminates.

The phases for MODIFY RECOVERY are:

© Copyright IBM Corp. 1983, 2001

UTILINIT	Initialization and setup
MODIFY	Deleting records
UTILTERM	Cleanup.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

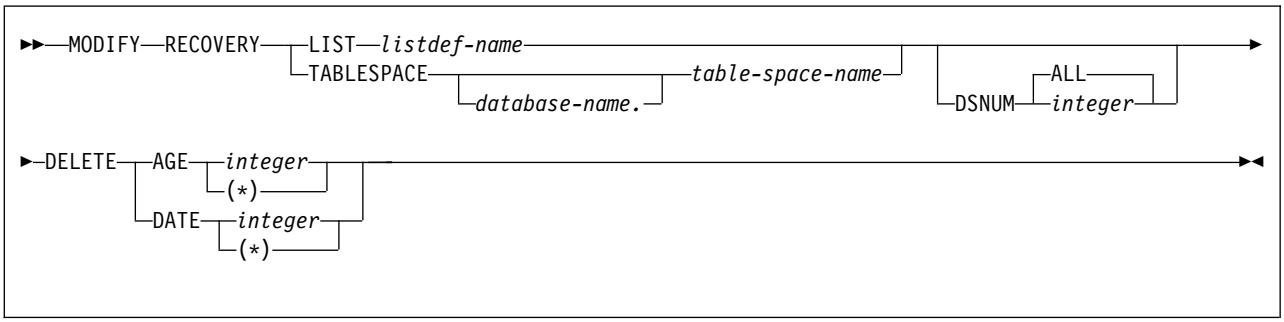

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name that contains only table spaces. You can specify one LIST keyword per MODIFY RECOVERY control statement. Do not specify LIST with the TABLESPACE keyword. Utility list manager will invoke MODIFY once for each table space in the list. For more information about LISTDEF specifications, see "Generating object lists using LISTDEF" on page 144.

TABLESPACE *database-name.table-space-name*
Specifies the database and the table space for which records are to be deleted.

database-name
Specifies the name of the database to which the table space belongs. *database-name* is optional.

The **default** is **DSNDB04**.

table-space-name
Specifies the name of the table space.

DSNUM *integer*
Identifies a single partition or data set of the table space for which records are to be deleted; ALL deletes records for the entire data set and table space.

integer is the number of a partition or data set.

The **default** is **ALL**.

For a partitioned table space, *integer* is its partition number. The maximum is 254.

For a nonpartitioned table space, use the data set integer at the end of the data set name as cataloged in the VSAM catalog. If image copies are taken by partition or data set and you specify DSNUM ALL, then the table space is placed in COPY-pending status if a full image copy of the entire table space does not exist. The data set name has this format:

catname.DSNDB*x*.*dbname*.*tsname*.*y*0001.A*nnn*

where: *y* is either I or J and *nnn* is the data set integer.

If you specify DSNUM, MODIFY RECOVERY does not delete any SYSCOPY records for the partition that have an RBA greater than that of the earliest point to which the entire table space could be recovered. That point might indicate a full image copy, a LOAD operation with LOG YES or a REORG operation with LOG YES.

DELETE
Indicates that records are to be deleted. See the DSNUM description above for restrictions on deleting partition statistics.

AGE *integer*
Deletes all SYSCOPY records older than a specified number of days.

integer is the number of days, and can range from 0 to 32767. Records created today are of age 0, and cannot be deleted by this option.

(*) deletes all records, regardless of their age.

DATE *integer*
Deletes all records written before a specified date.

integer may use either an 8 or 6 character format. You must specify a year (*yyyy* or *yy*), month (*mm*), and day (*dd*) in the form *yyyymmdd* or *yymmdd*. DB2 processing queries the system clock and converts 6-character dates to the most recent, previous 8-character equivalent.

(*) deletes all records, regardless of the date on which they were written.

Instructions for running MODIFY RECOVERY

To run MODIFY RECOVERY you must:

1. Read “Before running MODIFY RECOVERY” on page 234 in this chapter.
2. Prepare the necessary data sets, as described in “Data sets used by MODIFY RECOVERY” on page 234.
3. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for MODIFY RECOVERY, see “Sample control statements” on page 236.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 234. (For a complete description of the syntax and options for MODIFY RECOVERY, see “Syntax diagram” on page 232.)
5. Check the compatibility table in “Concurrency and compatibility” on page 235 if you want to run other jobs concurrently on the same target objects.
6. You can restart a MODIFY RECOVERY utility job, but it starts from the beginning again. You can terminate MODIFY RECOVERY with the TERM

UTILITY command. For guidance in restarting online utilities, see "Terminating or restarting MODIFY RECOVERY" on page 235.

7. Run MODIFY RECOVERY.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Before running MODIFY RECOVERY

Printing SYSCOPY records with REPORT RECOVERY: If you use MODIFY RECOVERY to delete SYSCOPY records, we recommend that you first use the REPORT utility to view all SYSCOPY records for the object at the specified site to avoid deleting the wrong records.

Removing RECOVER-pending status: You cannot run MODIFY RECOVERY on a table space that is in RECOVER-pending status. See "Chapter 21. RECOVER" on page 275 for information about resetting the RECOVER-pending status.

Data sets used by MODIFY RECOVERY

Table 41 describes the data sets used by MODIFY RECOVERY. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 41. Data sets used by MODIFY RECOVERY

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table space
Object for which records are to be deleted. It is named in the MODIFY RECOVERY control statement and is accessed through the DB2 catalog.

Creating the control statement

See "Syntax diagram" on page 232 for MODIFY RECOVERY syntax and option descriptions. See "Sample control statements" on page 236 for examples of MODIFY RECOVERY usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

"Deleting SYSLGRNX rows"
"Deleting all image copy entries" on page 235
"Deleting recovery rows for indexes" on page 235
"Reclaiming space in the DBD" on page 235
"Improving REORG performance after adding a column" on page 235

Deleting SYSLGRNX rows

If you take image copies at the data set level or partition level only, then you can use MODIFY RECOVERY with the DSNUM option to delete SYSLGRNX rows for the data set or partition. For partitioned table spaces, MODIFY RECOVERY deletes partition-level SYSLGRNX records only if there are no image copies of the entire

table space. Otherwise, if full copies exist for the entire table space and you specify DSNUM, MODIFY RECOVERY returns a message saying that no records were deleted.

If you omit DSNUM or specify DSNUM ALL, then MODIFY RECOVERY deletes all SYSLGRNX records pertaining to the entire table space and individual data sets and partitions.

Deleting all image copy entries

MODIFY RECOVERY allows you to delete all image copy entries for a table space or data set. In this case MODIFY RECOVERY:

- Issues message DSNU572I.
- Sets the COPY-pending restriction.
- Gives return code 4.

Deleting recovery rows for indexes

When you perform MODIFY RECOVERY on a table space, utility processing deletes SYSIBM.SYSCOPY and SYSIBM.SYSLGRNX rows that meet the AGE and DATE criteria for related indexes that were defined with COPY YES.

Reclaiming space in the DBD

To reclaim space in the DBD when you drop a table, use the following procedure:

1. Commit the drop.
2. Run the REORG utility.
3. Run the COPY utility to make a full image copy of the table space.
4. Run MODIFY RECOVERY with the DELETE option to delete all previous image copies.

Improving REORG performance after adding a column

After you add a column to a table space, the next REORG of the table space materializes default values for the added column by decompressing all rows of the table space during the UNLOAD phase and then compressing them again during the RELOAD phase. Subsequently, each REORG job for the table space repeats this processing in the UNLOAD and RELOAD phases. Use the following procedure to avoid repeating the compression cycle with each REORG:

1. Run the REORG utility on the table space.
2. Run the COPY utility to make a full image copy of the table space.
3. Run MODIFY RECOVERY with the DELETE option to delete all previous image copies. MODIFY RECOVERY changes the alter added column status only if there are SYSCOPY rows to delete.

Terminating or restarting MODIFY RECOVERY

MODIFY RECOVERY can be terminated in any phase without any integrity exposure.

You are permitted to restart the MODIFY RECOVERY utility, but it starts from the beginning again.

For more guidance in restarting online utilities, see “Restarting an online utility” on page 46.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Table 42 shows the restrictive state the utility sets on the target object.

Table 42. Claim classes of MODIFY RECOVERY operations. Use of claims and drains; restrictive states set on the target object.

Target	MODIFY RECOVERY
Table space or partition	UTRW

Legend:

UTRW - Utility restrictive state - Read/Write access allowed.

MODIFY RECOVERY can run concurrently on the same target object with any utility **except the following**:

- COPY TABLESPACE
- LOAD
- MERGECOPY
- MODIFY RECOVERY
- RECOVER TABLESPACE
- REORG TABLESPACE

The target object can be a table space or partition.

Sample control statements

Example 1: Delete SYSCOPY records by age. For the table space containing the employee table, delete all SYSCOPY records older than 90 days.

```
//STEP1    EXEC DSNUPROC,UID='IUIQU2UD.MODRCV1',
//            UTPROC='',SYSTEM='DSN'
//SYSIN    DD *
MODIFY RECOVERY TABLESPACE DSN8D71A.DSN8S71E DELETE AGE(90)
/*
```

Example 2: Delete SYSCOPY records by date. For the table space containing the department table, delete all SYSCOPY records written before 10 September 1998.

```
MODIFY RECOVERY TABLESPACE DSN8D71A.DSN8S71D
  DELETE DATE(19980910)
```

Chapter 17. MODIFY STATISTICS

The MODIFY STATISTICS online utility deletes unwanted statistics history records from the corresponding catalog tables. You can remove statistics history records that were written before a specific date, or you can remove records of a specific age. You can delete records for an entire table space, index space, or index.

Run MODIFY STATISTICS regularly to clear outdated information from the statistics history catalog tables. By deleting outdated information from those tables, you can improve performance for processes that access data from those tables.

For a diagram of MODIFY STATISTICS syntax and a description of available options, see "Syntax and options of the control statement". For detailed guidance on running this utility, see "Instructions for running MODIFY STATISTICS" on page 239.

Output: MODIFY STATISTICS deletes rows from the following catalog tables:
- SYSIBM.SYSCOLDIST_HIST
- SYSIBM.SYSCOLUMNS_HIST
- SYSIBM.SYSINDEXES_HIST
- SYSIBM.SYSINDEXPART_HIST
- SYSIBM.SYSINDEXSTATS_HIST
- SYSIBM.SYSLOBSTATS_HIST
- SYSIBM.SYSTABLEPART_HIST
- SYSIBM.SYSTABSTATS_HIST
- SYSIBM.SYSTABLES_HIST

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority.

A user ID with installation SYSOPR authority can also execute MODIFY STATISTICS, but only on a table space in the DSNDB01 or DSNDB06 database.

Execution phases of MODIFY STATISTICS:The MODIFY STATISTICS utility operates in these phases:

The phases for MODIFY STATISTICS are:

Phase	Description
UTILINIT	Initialization and setup
MODIFYS	Deleting records
UTILTERM	Cleanup.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

© Copyright IBM Corp. 1983, 2001

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

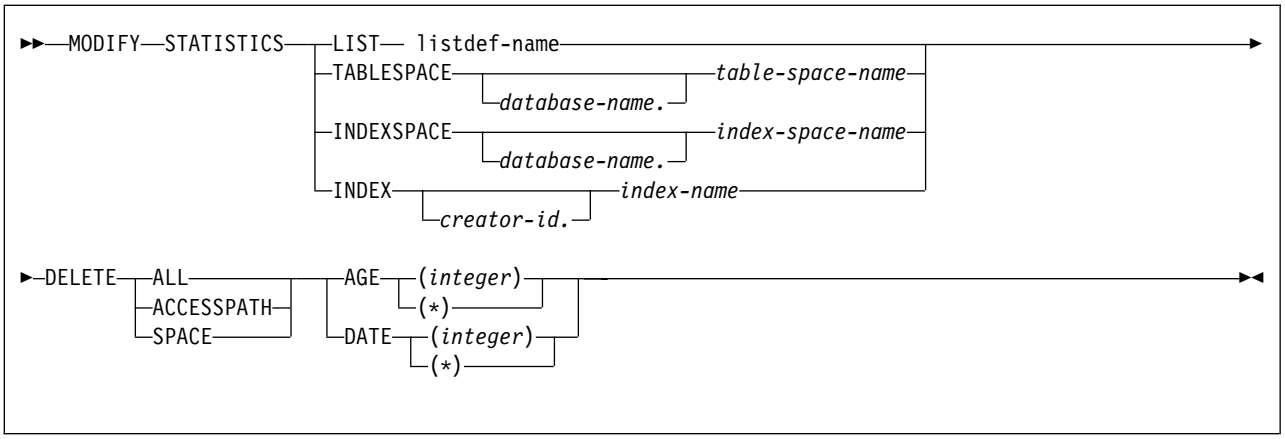

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

LIST*listdef-name*
Specifies the name of a previously-defined LISTDEF list name. You cannot repeat the LIST keyword or specify it with TABLESPACE, INDEXSPACE, or INDEX.

The list can contain index spaces, table spaces or a mixture of both. Utility list manager invokes MODIFY STATISTICS once per object in the list.

TABLESPACE *database-name.table-space-name*
Specifies the database and the table space for which catalog history records are to be deleted.

database-name
Specifies the name of the database to which the table space belongs. *database-name* is optional.

The **default** is **DSNDB04**.

table-space-name
Specifies the name of the table space for which statistics are to be deleted.

INDEXSPACE *database-name.index-space-name*
Specifies the qualified name of the index space for which catalog history information is to be deleted; the name listed in the SYSIBM.SYSINDEXES table.

database-name
Optionally specifies the name of the database to which the index space belongs.

The **default** is **DSNDB04**.

index-space-name
Specifies the name of the index space for which the statistics are to be deleted.

INDEX *creator-id.index-name*
Specifies the index for which catalog history information is to be deleted.

creator-id
Optionally specifies the creator of the index.

The **default** is **DSNDB04**.

index-name
Specifies the name of the index for which the statistics are to be deleted.

DELETE
Indicates that records are to be deleted.

ALL Deletes all statistics history rows that are related to the specified object from all catalog history tables.

ACCESSPATH
Deletes all access-path statistics history rows that are related to the specified object from the following history tables:

- SYSIBM.SYSINDEXPART_HIST
- SYSIBM.SYSTABLEPART_HIST
- SYSIBM.SYSLOBSTATS_HIST

SPACE
Deletes all space-tuning statistics history rows related to the specified object from the following history tables:

- SYSIBM.SYSINDEXPART_HIST
- SYSIBM.SYSTABLEPART_HIST
- SYSIBM.SYSLOBSTATS_HIST

AGE (*integer***)**
Deletes all statistics history rows, related to the specified object, that are older than a specified number of days.

(*integer***)**
Specifies the number of days in a range from 0 to 32 767. This option cannot delete records that are created today (age 0).

(*) Deletes all records, regardless of their age.

DATE (*integer***)**
Deletes all records written before a specified date.

(*integer***)**
Specifies the date in an eight character format. Specify a year (*yyyy*), month (*mm*), and day (*dd*) in the form *yyyymmdd*.

(*) Deletes all records, regardless of the date on which they were written.

Instructions for running MODIFY STATISTICS

To run MODIFY STATISTICS you must:

1. Prepare the necessary data sets, as described in “Data sets used by MODIFY STATISTICS” on page 240.
2. Create JCL statements by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for MODIFY STATISTICS, see “Sample control statements” on page 241.)

3. Prepare a utility control statement by specifying the options for the tasks that you want to perform, as described in "Instructions for specific tasks". (For a complete description of the syntax and options for MODIFY STATISTICS, see "Syntax diagram" on page 238.)
4. Check the compatibility table in "Concurrency and compatibility" on page 241 if you want to run other jobs concurrently on the same target objects.
5. Restart a MODIFY STATISTICS utility job (it starts from the beginning again) or terminate MODIFY STATISTICS by using the TERM UTILITY command. For guidance in restarting online utilities, see "Restarting an online utility" on page 46.
6. Run MODIFY STATISTICS.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of methods you can use to execute DB2 utilities.

Data sets used by MODIFY STATISTICS

Table 43 describes the data sets used by MODIFY STATISTICS. Include statements in your JCL for each required data set and any optional data sets you want to use.

Table 43. Data sets used by MODIFY STATISTICS

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement	Yes
SYSPRINT	Output data set for messages	Yes

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table space or index space
Object for which records are to be deleted. This object is named in the MODIFY STATISTICS control statement and is accessed through the DB2 catalog.

Creating the control statement

See "Syntax diagram" on page 238 for MODIFY STATISTICS syntax and option descriptions. See "Sample control statements" on page 241 for examples of MODIFY STATISTICS usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement.

Deciding which statistics history rows to delete

After analyzing trends by using the relevant historical catalog information and possibly taking actions based on this information, consider deleting all or part of the statistics history catalog rows. Deleting outdated information from the statistics history catalog tables can improve performance for processes that access data from those tables, and you also free up space in the catalog. Then, the next time you update the relevant statistics by using RUNSTATS TABLESPACE, REBUILD INDEX, or REORG INDEX, DB2 repopulates the statistics history catalog tables with more recent historical data. Examining this data lets you determine the efficacy of any adjustments you made as a result of your previous analysis.

Be aware that when you manually insert, update or delete catalog information, DB2 does not store the historical information for those operations in the historical catalog tables.

Deleting specific statistics history rows

MODIFY STATISTICS lets you delete all or portions of statistics history rows for a table space, an index space or an index.

You can choose to delete only the statistics rows that relate to access path selection by specifying the ACCESSPATH option, or to delete the rows that relate to space statistics by using the SPACE option. To delete rows in the all statistics history catalog tables, including the SYSIBM.SYSTABLES_HIST catalog table, you must specify the DELETE ALL option in the utility control statement.

You can also delete rows that meet the AGE and DATE criteria by specifying the corresponding keywords for a particular object.

Terminating or restarting MODIFY STATISTICS

You can terminate the MODIFY STATISTICS utility in any phase.

You can restart the MODIFY STATISTICS utility, but it starts from the beginning again.

For more guidance in restarting online utilities, see “Restarting an online utility” on page 46.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities that operate on different partitions of the same table space or index space are compatible.

Table 44 shows the restrictive state that the utility sets on the target object.

Table 44. Claim classes of MODIFY STATISTICS operations. Use of claims and drains; restrictive states set on the target object.

Target	MODIFY STATISTICS
Table space, index, or index space	UTRW

Legend:

UTRW - Utility restrictive state - Read/Write access allowed.

The target object can be a table space, an index, or an index space.

Sample control statements

Example 1: Delete SYSIBM.SYSTABLES_HIST records by age. For the table space containing the employee table, delete all statistics history records older than 60 days.

```
//STEP1    EXEC DSNUPROC,UID='IUIQU2UD.MODSTAT1',
//            UTPROC='',SYSTEM='V71A'
//SYSIN    DD *
```

```
  MODIFY STATISTICS TABLESPACE DSN8D71A.DSN8S71E
    DELETE ALL
    AGE 60
/*
```

Example 2: Delete accesspath records for a list by date. Define a list named M1. Delete all accesspath statistics history records for list M1 that are older than April 17, 2000.

```
//STEP9    EXEC DSNUPROC,UID='JUOEU115.MDFYL9',
//         UTPROC='',
//         SYSTEM='SSTR'
//SYSUT1   DD DSN=JUOEU115.MDFYL.STEP9.SYSUT1,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(8000,(200,20),,,ROUND)
//SYSERR   DD DSN=JUOEU115.MDFYL.STEP9.SYSERR,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SORTOUT  DD DSN=JUOEU115.MDFYL.STEP9.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SYSIN    DD *

   LISTDEF M1 INCLUDE TABLESPACE DBOE1501.TLOE1501
   MODIFY STATISTICS LIST M1
                     DELETE ACCESSPATH DATE(20000417)
/*
```

Example 3: Delete space-tuning statistic records for an index by age. Delete all space-tuning statistic records for index ADMF001.IXOE15S1 that are older than one day.

```
//STEP9    EXEC DSNUPROC,UID='JUOEU115.MOFYS9',
//         UTPROC='',
//         SYSTEM='SSTR'
//SYSUT1   DD DSN=JUOEU115.MOFYS.STEP9.SYSUT1,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(8000,(200,20),,,ROUND)
//SYSERR   DD DSN=JUOEU115.MOFYS.STEP9.SYSERR,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SORTOUT  DD DSN=JUOEU115.MOFYS.STEP9.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SYSIN    DD *

   MODIFY STATISTICS INDEX ADMF001.IXOE15S1
                     DELETE SPACE AGE 1
/*
```

Example 4: Delete all statistic history records for an index space regardless of date. For index space DBOE1501.IUOE1501, delete all statistics history records.

```
//STEP8    EXEC DSNUPROC,UID='JUOEU115.MDFYL8',
//         UTPROC='',
//         SYSTEM='SSTR'
//SYSUT1   DD DSN=JUOEU115.MDFYL.STEP8.SYSUT1,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(8000,(200,20),,,ROUND)
//SYSERR   DD DSN=JUOEU115.MDFYL.STEP8.SYSERR,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
//SORTOUT  DD DSN=JUOEU115.MDFYL.STEP8.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(6000,(20,20),,,ROUND)
```

```
//SYSIN    DD *
   MODIFY STATISTICS INDEXSPACE DBOE1501.IUOE1501
                     DELETE ALL DATE (*)
/*
```

Chapter 18. OPTIONS

The OPTIONS utility control statement specifies processing options that are applicable across many utility executions in a job step. By specifying various options you can:

- Preview utility control statements
- Preview LISTDEF or TEMPLATE definitions
- Override library names for LISTDEF lists or TEMPLATEs
- Specify how to handle errors during list processing
- Alter the return code for warning messages
- Restore all default options

See “Syntax and options of the control statement” for details.

Output: The OPTIONS control statement sets the processing options specified for the duration of the job step, or until replaced by another OPTIONS control statement within the same job step.

Authorization required: The OPTIONS control statement performs setup for subsequent control statements. The OPTIONS statement itself requires no privileges to execute.

Execution phases of OPTIONS:The OPTIONS control statement executes entirely in the UTILINIT phase, performing setup for the subsequent utility.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

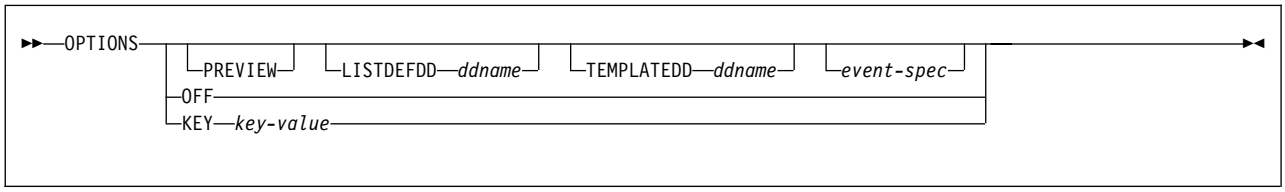

© Copyright IBM Corp. 1983, 2001

OPTIONS

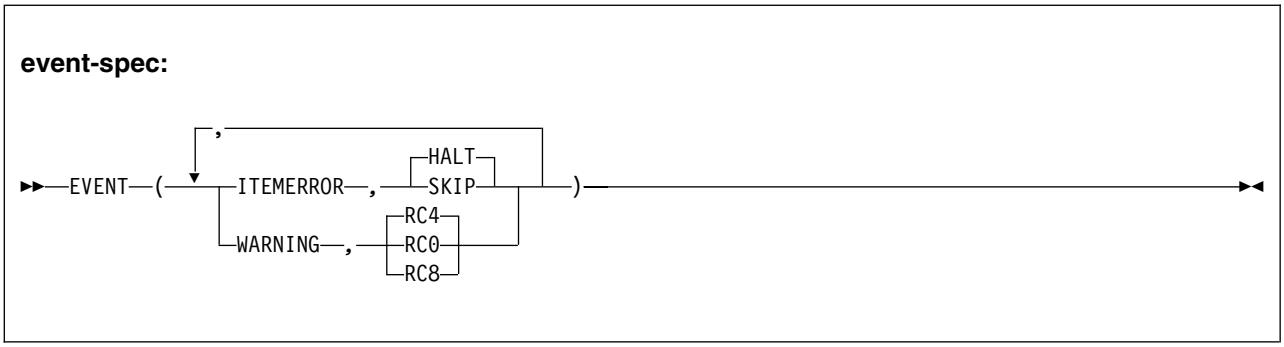

Option descriptions

For a description of parsing utility statements, and interpreting option identifiers and specifications, see “Control statement coding rules” on page 23.

PREVIEW Executes in PREVIEW mode the utility control statements that follow. It checks for syntax errors in all utility control statements, but normal utility execution will not take place. If syntax is valid, the utility will expand all LISTDEF lists and TEMPLATE DSNs that appear in SYSIN and print results to the SYSPRINT data set.

PREVIEW will evaluate and expand all LISTDEFs into an actual list of table spaces or index spaces. It evaluates TEMPLATE DSNs into actual data set names through variable substitution. It also expands lists from the SYSLISTD DD and TEMPLATE DSNs from the SYSTEMPL DD that is referenced by a utility invocation.

An alternative to the OPTIONS PREVIEW statement is the PREVIEW JCL PARM. These two functions are identical. JCL PARM method is preferred for previewing an existing set of utility control statements. The OPTION PREVIEW control statement method is preferred when invoking DB2 utilities via a stored procedure.

A new value will be allowed for the third JCL PARM of DSNUTILB and for the &UTPROC. variable of DSNUPROC:

```
//STEP1    EXEC DSNUPROC,UID='JULTU106.RECOVE1',
//         UTPROC='PREVIEW',SYSTEM='SSTR'
```

The PARM value PREVIEW will cause the utility control statements in that job step to be processed for preview only. The LISTDEF and TEMPLATE control statements will be expanded but no utility execution will be performed.

OPTIONS PREVIEW is identical to the PREVIEW JCL parameter, except that the preview for OPTIONS PREVIEW may be turned off in the same job step by a later OPTIONS statement. Absence of the PREVIEW keyword in the OPTION control statement turns off preview processing but it does not override the PREVIEW JCL parameter which, if specified, remains in affect for the entire job step.

LISTDEFDD *ddname*
Specifies the *ddname* of the LISTDEF definition library. The **default**

value is **SYSLISTD**. A LISTDEF library is a data set that contains only LISTDEF control statements. This data set is processed only when a referenced LIST is not found in SYSIN. See LISTDEF CONTROL STATEMENT for details.

TEMPLATEDD *ddname*
Specifies the *ddname* of the TEMPLATE definition library. The **default** value is **SYSTEMPL**. A TEMPLATE library is a data set that contains only TEMPLATE control statements. This data set is processed only when a referenced name does not exist in the job step as a DD name and is not found in SYSIN as a TEMPLATE name. See TEMPLATE CONTROL STATEMENT for details.

EVENT Specifies one or more pairs of utility processing events and the matching action for the event. Not all actions are valid for all events.

OFF Restores all default options. OPTIONS OFF does not override the PREVIEW JCL parameter which, if specified, remains in effect for the entire job step. No other OPTIONS keywords may be specified with OPTIONS OFF.

OPTIONS OFF is equivalent to OPTIONS LISTDEFDD SYSLISTD TEMPLATEDD SYSTEMPL EVENT (ITEMERROR, HALT, WARNING, RC4).

KEY Only use OPTIONS KEY when instructed by IBM Service. OPTIONS KEY is followed by a single operand which will be provided when needed, by IBM service.

ITEMERROR

Specifies how utility processing is to handle errors during list processing. Specifically, if an error message, with return code 8 is received, should processing halt, or skip the error and continue to the end of the control statement. By default, utility processing will HALT. The ITEMERROR event does not include abnormal terminations (ABENDS).

Note that for the QUIESCE utility, the indexes for the table spaces in the list, if any, are considered as list items for the purposes of the ITEMERROR event. ITEMERROR will affect how errors are handled on both the table spaces and the indexes.

HALT
Specifies that the utility is to HALT after the event.

SKIP
Ignores the event and skips the list item. Processing continues with the next item in the list.

SKIP processing only applies during the processing of a valid list. If a utility detects that a LIST is invalid for the utility being invoked, it will be rejected with an error message and the processing of the list will not be initiated. SKIP processing does not apply in this case.

If any of the items in a list is skipped, the utility produces a return code of 8 which terminates the job step. In this example, use the following coding:

```
SYSIN DD *
      OPTIONS ITEMERROR, SKIP
      COPY LISTA
      COPY LISTB
```

If LISTA contains ten objects and one object produces RC8 during the COPY, the other nine objects in the list are copied successfully. The job step ends with RC8 and COPY LISTB is not executed.

WARNING Specifies a warning message event.

Use WARNING to alter the return code for warning messages. The return code shown by message DSNU010I may be altered by this option. If it is altered, message DSNU1024I will be issued to document the new return code.

Action choices are RC0, RC4, or RC8.

RC0
Lowers the final return code of a single utility invocation that ends in a return code 4 to a return code of 0. Use RC0 to force a return code of 0 for warning messages.

This option should generally be avoided. Warning messages document situations which may or may not be as expected and should not be disregarded. Use this option only when return code 4 is expected, is acceptable, and other mechanisms are in place to validate the results of a utility execution.

RC4
Specifies that return codes for warning messages will remain unchanged. Use RC4 to override a previous OPTIONS WARNING specification in the same job step.

RC8
Raises the final return code of a single utility invocation that ends in a return code 4 to a return code of 8. Use RC8 to force a return code of 8 for warning messages. The return code of 8 will cause the job step to terminate and subsequent utility control statements will not be executed.

Instructions for using OPTIONS

To execute utility control statements in preview mode, use OPTIONS PREVIEW. Control statements will be previewed for use with LISTDEF lists and TEMPLATEs but the options specified are not actually executed.

To use different LISTDEF and TEMPLATE libraries, specify OPTIONS LISTDEFDD and OPTIONS TEMPLATEDD to override the names of the optional library data sets.

To alter settings for warning return codes and error handling during list processing, use OPTIONS EVENT to override the standard utility processing behaviors.

You can repeat an OPTIONS control statement within the SYSIN DD card. If you repeat the control statement, it entirely replaces any prior OPTIONS control statement.

Concurrency and compatibility

OPTIONS is a utility control statement used to setup an environment for another utility to follow. Until the OPTIONS is referenced by a specific utility it is simply stored for later reference. When referenced by another utility, the list will be expanded. At that time, the concurrency and compatibility restrictions of that utility apply, with the additional restriction that the catalog tables necessary to expand the list must be available for read only access.

Sample control statements

Example 1: Preview a COPY utility using LISTDEF and TEMPLATE definition. Preview a COPY utility using a LISTDEF definition from the PAYLIST data set and TEMPLATE definitions from the PAYTMPL data set.

```
OPTIONS PREVIEW
TEMPLATE COPYLOC UNIT(SYSDA)
         DSN(&DB..&TS..D&JDATE..&STEPNAME..COPY&IC.&LOCREM.&PB.)
         DISP(NEW,CATLG,CATLG) SPACE(200,20) TRK
         VOLUMES(SCR03)
TEMPLATE COPYREM UNIT(SYSDA)
         DSN(&DB..&TS..&UT..T&TIME..COPY&IC.&LOCREM.&PB.)
         DISP(NEW,CATLG,CATLG) SPACE(100,10) TRK
LISTDEF CPYLIST INCLUDE TABLESPACES DATABASE DBLT0701
COPY LIST CPYLIST FULL YES
     COPYDDN(COPYLOC,COPYLOC)
     RECOVERYDDN(COPYREM,COPYREM)
     SHRLEVEL REFERENCE
```

Example 2: Using OPTIONS with COPY. Assume that two versions of an application exist, using two different sets of database objects and that LISTDEFs and TEMPLATEs exist for both versions in two different sets of library data sets. COPY both sets of database objects.

```
OPTIONS LISTDEFDD V1LIST TEMPLATEDD V1TEMPL
COPY LIST PAYTBSP COPYDDN(PAYTEMP1,PAYTEMP1)
OPTIONS LISTDEFDD V2LIST TEMPLATEDD V2TEMPL
COPY LIST PAYTBSP COPYDDN(PAYTEMP1,PAYTEMP1)
```

Example 3: Using OPTIONS with MODIFY RECOVERY. Assume that a MODIFY RECOVERY utility will be deleting all image copies. This is expected and acceptable. The return code 4 is to be ignored. Before the next utility, turn this option off.

```
OPTIONS EVENT(WARNING,RC0)
MODIFY RECOVERY TABLESPACE A.B DELETE AGE(*)
OPTIONS OFF
MODIFY RECOVERY TABLESPACE C.D DELETE AGE(30)
```

Example 4: Preview COPY utilities, then override the preview with ITEMERROR and SKIP options.

```
   OPTIONS PREVIEW
   LISTDEF COPY1_LISTDEF
               INCLUDE TABLESPACES TABLESPACE DSNDB01.SPT01
               INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSDBASE
               INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSDBAUT
               INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSGPAUT
               INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSGROUP
               INCLUDE TABLESPACES TABLESPACE DBA91302.T?A9132*
   LISTDEF COPY2_LISTDEF
               INCLUDE TABLESPACES TABLESPACE DBA91303.TLA9133A
               INCLUDE TABLESPACES TABLESPACE DBA91303.TSA9133B
```

```
|                                    INCLUDE TABLESPACES TABLESPACE DBA91303.TPA9133C
|                                    INCLUDE TABLESPACES TABLESPACE DBA91304.TLA9134A
|                                    INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSUSER
|                                    INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSVIEWS
|                                    INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSSTATS
|                                    INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSDDF
|                                    INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSOBJ
|                               LISTDEF COPY3_LISTDEF
|                                    INCLUDE TABLESPACES TABLESPACE DBA91304.TSA9134B
|                                    INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSHIST
|                                    INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSGRTNS
|                                    INCLUDE TABLESPACES TABLESPACE DSNDB06.SYSJAVA
|                                    INCLUDE TABLESPACES TABLESPACE DBA91304.TPA9134C
|                               OPTIONS EVENT(ITEMERROR,SKIP)
|                               TEMPLATE TMP1 UNIT(SYSDA) DISP(MOD,CATLG,CATLG)
|                                        VOLUMES(SCR03)
|                                        DSN(DH109013.&TS..COPY&ICTYPE.&LOCREM.&PRIBAC.)
|                               COPY LIST COPY1_LISTDEF SHRLEVEL REFERENCE
|                                    COPYDDN (TMP1)
|                                    RECOVERYDDN (TMP1)
|                                    FULL YES
|                               COPY LIST COPY2_LISTDEF SHRLEVEL REFERENCE
|                                    COPYDDN (TMP1,TMP1)
|                                    FULL YES
|                               COPY LIST COPY3_LISTDEF SHRLEVEL REFERENCE
|                                    COPYDDN (TMP1,TMP1)
|                                    RECOVERYDDN (TMP1,TMP1)
|                                    FULL YES
```

Chapter 19. QUIESCE

The QUIESCE online utility establishes a quiesce point (the current log RBA or log record sequence number (LRSN)) for a table space, partition, table space set, or list of table spaces and table space sets, and records it in the SYSIBM.SYSCOPY catalog table. A successful QUIESCE improves the probability of a successful RECOVER or COPY. You should run QUIESCE frequently between regular executions of COPY to establish regular recovery points for future point in time recovery.

For a diagram of QUIESCE syntax and a description of available options, see "Syntax and options of the control statement". For detailed guidance on running this utility, see "Instructions for running QUIESCE" on page 253.

Output: QUIESCE writes changed pages from the table spaces to disk. The catalog table SYSIBM.SYSCOPY records the current RBA and the timestamp of the quiesce point. A row with ICTYPE='Q' is inserted into SYSCOPY for each table space quiesced. DB2 also inserts a SYSCOPY row with ICTYPE='Q' for any indexes (defined with the COPY YES attribute) over a table space being quiesced. (Table spaces DSNDB06.SYSCOPY, DSNDB01.DBD01, and DSNDB01.SYSUTILX are an exception: their information is written to the log.)

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- IMAGCOPY privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority.

An ID with installation SYSOPR authority can also execute QUIESCE, but only on a table space in the DSNDB01 or DSNDB06 database.

Execution phases of QUIESCE: The QUIESCE utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup
QUIESCE	Determining the quiesce point and updating the catalog
UTILTERM	Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

© Copyright IBM Corp. 1983, 2001

QUIESCE

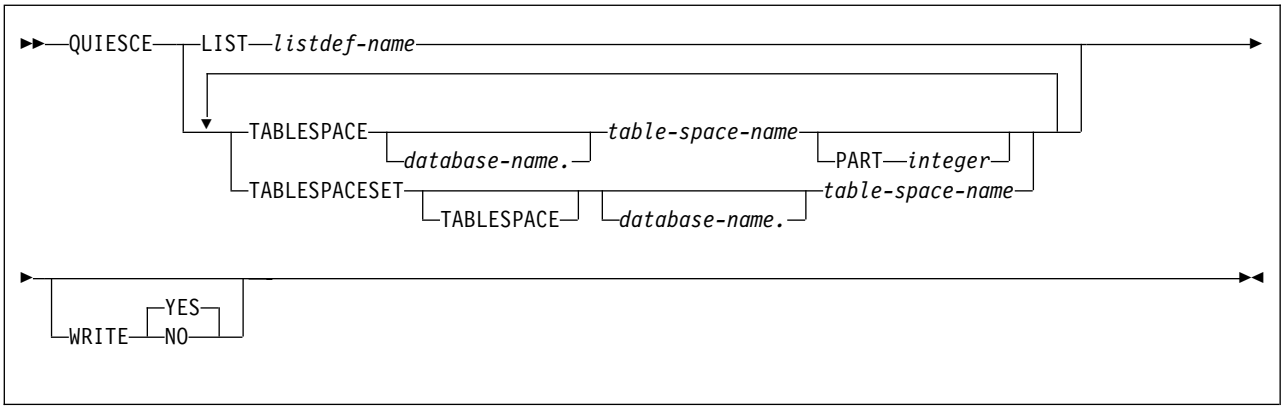

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name that contains only table spaces. The utility allows one LIST keyword for each QUIESCE control statement. Do not specify LIST with the TABLESPACE or TABLESPACESET keyword. Utility list manager will invoke QUIESCE once for the entire list. For the QUIESCE utility, the related index spaces are considered to be list items for the purposes of OPTIONS ITEMERROR processing. Therefore, the utility behavior during processing of related indexes can be altered with the OPTIONS ITEMERROR statement. For more information about LISTDEF specifications, see “Generating object lists using LISTDEF” on page 144.

TABLESPACE *database-name.table-space-name*
For QUIESCE TABLESPACE, specifies the table space (and, optionally, the database to which it belongs) being quiesced.

For QUIESCE TABLESPACESET, specifies a table space (and, optionally, the database to which it belongs) in the table space set being quiesced.

database-name
Optionally specifies the name of the database to which the table space belongs. The **default** is **DSNDB04**.

table-space-name
Specifies the name of the table space to be quiesced. You can specify DSNDB01.SYSUTILX, but it cannot be included in a list with other table spaces to be quiesced.

PART *integer*
Identifies a partition to be quiesced.

integer is the number of the partition and must be in the range from 1 to the number of partitions defined for the table space. The maximum is 254.

TABLESPACESET
Specifies that all of the referentially related table spaces in the table space set are to be quiesced. For the purposes of the QUIESCE utility, a table space set is either:

- A group of table spaces that have a referential relationship.
- A base table space with all of its LOB table spaces.

TABLESPACE *database-name.table-space-name*
Specifies the table space name (and, optionally, the database to which it belongs) being quiesced. For QUIESCE TABLESPACESET, the TABLESPACE keyword is optional.

WRITE
Specifies whether to write the changed pages from the table spaces and index spaces to disk.

YES
Establishes a quiesce point and writes the changed pages from the table spaces and index spaces to disk.

The **default** is **YES**.

NO
Establishes a quiesce point but does not write the changed pages from the table spaces and index spaces to disk.

Instructions for running QUIESCE

To run QUIESCE, you must:

1. Read "Before running QUIESCE" in this chapter.
2. Prepare the necessary data sets, as described in "Data sets used by QUIESCE".
3. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for QUIESCE, see "Sample control statements" on page 258.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 254. (For a complete description of the syntax and options for QUIESCE, see "Syntax and options of the control statement" on page 251.)
5. Check the compatibility table in "Concurrency and compatibility" on page 256 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the QUIESCE job doesn't complete, as described in "Terminating or restarting QUIESCE" on page 256.
7. Run QUIESCE.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Before running QUIESCE

You cannot run QUIESCE on a table space that is in COPY-pending, CHECK pending, RECOVER-pending, or auxiliary CHECK-pending status. See "Resetting the COPY-pending status" on page 209, "Resetting CHECK-pending status" on page 63, "Resetting the REBUILD-pending status" on page 209, and "Appendix C. Resetting an advisory or restrictive status" on page 677 for information about resetting these statuses.

Data sets used by QUIESCE

Table 45 on page 254 describes the data sets used by QUIESCE. Include statements in your JCL for each required data set, and any optional data sets you

want to use.

Table 45. Data sets used by QUIESCE

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table space
Object to be quiesced. It is named in the QUIESCE control statement and is accessed through the DB2 catalog. (If you want to quiesce only one partition of a table space, you must use the PART option in the control statement.)

Creating the control statement

See "Syntax and options of the control statement" on page 251 for QUIESCE syntax and option descriptions. See "Sample control statements" on page 258 for examples of QUIESCE usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

"Using QUIESCE for recovery"
"Using QUIESCE on catalog and directory objects"
"Obtaining a common quiesce point"
"Specifying a list of table spaces and table space sets" on page 255

Using QUIESCE for recovery

You can recover a table space to its quiesce point with the RECOVER TABLESPACE utility. See "Chapter 21. RECOVER" on page 275 for information about the RECOVER TABLESPACE utility.

Using QUIESCE on catalog and directory objects

You can quiesce DSNDB01.SYSUTILX, but DSNDB01.SYSUTILX must be the only table space in the QUIESCE control statement.

If a point-in-time recovery is planned for the catalog and directory, a separate QUIESCE control statement for DSNDB06.SYSCOPY is required after you quiesce the other catalog and directory table spaces. ***Recommendation:*** Recover all catalog and directory objects to the current state.

Obtaining a common quiesce point

Use the QUIESCE TABLESPACESET utility to obtain a common quiesce point for table spaces that are related. For the purposes of the QUIESCE utility, a table space set is one or both of the following:

- A group of table spaces that have a referential relationship
- A base table space with all of its LOB table spaces

If you use QUIESCE TABLESPACE instead and do not include every member, you may encounter problems running RECOVER on the table spaces in the structure.

RECOVER checks if a complete table space set is recovered to a point in time. If the table space set is not complete, RECOVER places all dependent table spaces into CHECK-pending status.

You should QUIESCE and RECOVER the LOB table spaces to the same point in time as the associated base table space. A group of table spaces that have a referential relationship should all be quiesced to the same point in time.

When you QUIESCE WRITE YES a table space, each related index that was defined with COPY YES has a SYSIBM.SYSCOPY row inserted with ICTYPE='Q' to record the quiesce point.

Specifying a list of table spaces and table space sets

You can specify as many objects in your QUIESCE job as available memory in the batch address space and in the DSN1MSTR address space allows.

Be aware of the following considerations when you specify a list of objects to quiesce:

- Each table space set will be expanded into a list of table spaces that have a referential relationship, or a list containing a base table space with all of its LOB table spaces.
- If you specify a list of table spaces or table space sets to quiesce and duplicate a table space, utility processing will continue and the table space will only be quiesced once. QUIESCE will issue return code 4 and warning message DSNU533I to alert you of the duplication.
- If you specify the same table space twice in a list, using PART *n* in one specification, and PART *m* for the other specification, each partition is quiesced once.

Considerations for running QUIESCE

If a table space is in a pending status: If you run QUIESCE on a table space in COPY-pending, CHECK-pending, or RECOVER-pending status, it terminates as shown in Figure 15.

```
DSNU000I    DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = R92341Q
DSNU050I    DSNUGUTC - QUIESCE TABLESPACE UTQPD22A.UTQPS22D
                               TABLESPACE UTQPD22A.UTQPS22E
                               TABLESPACE UTQPD22A.EMPPROJA
DSNU471I  % DSNUQUIA COPY PENDING ON TABLESPACE UTQPD22A.EMPPROJA PROHIBITS PROCESSING
DSNU012I DSNUGBAC - UTILITY EXECUTION TERMINATED, HIGHEST RETURN CODE=8
```

Figure 15. QUIESCE and pending restrictions

If you run QUIESCE on a table space or an index space that is in a pending status, QUIESCE terminates and you might receive one or more of the messages shown in Figure 16 on page 256.

```
DSNU202I  CSECT  RECOVER PENDING ON TABLESPACE ... PROHIBITS PROCESSING
DSNU203I  CSECT  RECOVER PENDING ON INDEX ... PROHIBITS PROCESSING
DSNU204I  CSECT  PAGESET REBUILD PENDING ON INDEX ... PROHIBITS PROCESSING
DSNU208I  CSECT  GROUP BUFFER POOL RECOVER PENDING ON INDEX ... PROHIBITS PROCESSING
DSNU209I  CSECT  RESTART PENDING ON ... PROHIBITS PROCESSING
DSNU210I  CSECT  INFORMATIONAL COPY PENDING ON INDEX ... PROHIBITS PROCESSING
DSNU211I  CSECT  CHECK PENDING ON  ... PROHIBITS PROCESSING
DSNU214I  CSECT  REBUILD PENDING ON INDEX ... PROHIBITS PROCESSING
DSNU215I  CSECT  REFRESH PENDING ON  ... PROHIBITS PROCESSING
DSNU471I  CSECT  COPY PENDING ON TABLESPACE ... PROHIBITS PROCESSING
DSNU568I  CSECT  INDEX ... IS IN INFORMATIONAL COPY PENDING
```

Figure 16. Messages for pending restrictions on QUIESCE

If the write to disk fails: QUIESCE attempts to write pages of each table space to disk. If any of the following conditions is encountered, the write to disk fails:

- The table space has a write error range
- The table space has deferred restart pending
- An I/O error occurs

If any of the above conditions is true, QUIESCE will terminate with a return code of 4 and a DSNU473I warning message. In any case, the QUIESCE point is a valid point of consistency for recovery.

Terminating or restarting QUIESCE

If you use -TERM UTILITY to terminate QUIESCE when it is active, QUIESCE releases the drain locks on table spaces. If QUIESCE is stopped, the drain locks have already been released.

You can restart the QUIESCE utility, but it starts from the beginning again.

For more guidance in restarting online utilities, see “Restarting an online utility” on page 46.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Table 46 shows which claim classes QUIESCE drains and any restrictive state the utility sets on the target object.

Table 47 on page 257 shows which utilities can run concurrently with QUIESCE on the same target object. The target object can be a table space, an index space, or a partition of a table space or index space. If compatibility depends on particular options of a utility, that is also shown.

Table 46. Claim classes of QUIESCE operations. Use of claims and drains; restrictive states set on the target object.

Target	WRITE YES	WRITE NO
Table space or partition	DW/UTRO	DW/UTRO
Index or partition	DW/UTRO	
Nonpartitioning index	DW/UTRO	

Table 46. Claim classes of QUIESCE operations (continued). Use of claims and drains; restrictive states set on the target object.

Target	WRITE YES	WRITE NO

Legend:
- DW - Drain the write claim class - concurrent access for SQL readers
- UTRO - Utility restrictive state - read only access allowed

QUIESCE does not set a utility restrictive state if the target object is DSNDB01.SYSUTILX.

Table 47. QUIESCE compatibility

Action	QUIESCE
CHECK DATA	No
CHECK INDEX	Yes
CHECK LOB	Yes
COPY INDEXSPACE SHRLEVEL REFERENCE	Yes
COPY INDEXSPACE SHRLEVEL CHANGE	No
COPY TABLESPACE SHRLEVEL REFERENCE	Yes
COPY TABLESPACE SHRLEVEL CHANGE	No
DIAGNOSE	Yes
LOAD	No
MERGECOPY	Yes
MODIFY	Yes
QUIESCE	Yes
REBUILD INDEX	No
RECOVER INDEX	No
RECOVER TABLESPACE	No
REORG INDEX	No
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No
REORG TABLESPACE UNLOAD ONLY or EXTERNAL	Yes
REPAIR DUMP or VERIFY	Yes
REPAIR DELETE or REPLACE	No
REPORT	Yes
RUNSTATS	Yes
STOSPACE	Yes
UNLOAD	Yes

To run on DSNDB01.SYSUTILX, QUIESCE must be the only utility in the job step.

QUIESCE on SYSUTILX is an “exclusive” job; such a job can interrupt another job between job steps, possibly causing the interrupted job to time out.

Sample control statements

Example 1: Sample JCL for QUIESCE. Establish a quiesce point for three table spaces.

```
//STEP1    EXEC DSNUPROC,UID='IUIQU2UD.QUIESC2',
//            UTPROC='',SYSTEM='V71A'
//SYSIN    DD *
QUIESCE TABLESPACE DSN8D71A.DSN8S71D
        TABLESPACE DSN8D71A.DSN8S71E
        TABLESPACE DSN8D71A.DSN8S71P
//*
```

Example 2: Sample control statement for QUIESCE. Establish a quiesce point for the DSN8D71A.DSN8S71E and DSN8D71A.DSN8S71D table spaces.

```
QUIESCE TABLESPACE DSN8D71A.DSN8S71E TABLESPACE DSN8D71A.DSN8S71D
```

The following is output of the preceding command:

```
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TSLQ.STEP1
DSNU050I   DSNUGUTC - QUIESCE TABLESPACE DSN8D71A.DSN8S71E
                              TABLESPACE DSN8D71A.DSN8S71D
DSNU477I - DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.DSN8S71E
DSNU477I - DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.DSN8S71D
DSNU474I - DSNUQUIA - QUIESCE AT RBA 000000052708
DSNU475I   DSNUQUIB - QUIESCE UTILITY COMPLETE, ELAPSED TIME= 00:00:25
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

Example 3: QUIESCE point for a table space set. Establish a quiesce point for the table space set of the sample application.

```
QUIESCE TABLESPACESET TABLESPACE DSN8D71A.DSN8S71D
```

The following is output of the preceding command:

```
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TSLQ.STEP1
DSNU050I   DSNUGUTC - QUIESCE TABLESPACESET TABLESPACE DSN8D71A.DSN8S71D
DSNU477I - DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACESET DSN8D71A.DSN8S71D
DSNU477I - DSNUQUIA -    QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.DSN8S71D
DSNU477I - DSNUQUIA -    QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.DSN8S71E
DSNU477I - DSNUQUIA -    QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.PROJ
DSNU477I - DSNUQUIA -    QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.ACT
DSNU477I - DSNUQUIA -    QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.PROJACT
DSNU477I - DSNUQUIA -    QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.EMPPROJA
DSNU477I - DSNUQUIA -    QUIESCE SUCCESSFUL FOR TABLESPACE DSN8D71A.DSN8S1D
DSNU474I - DSNUQUIA - QUIESCE AT RBA 000000052708 AND AT LRSN 000000052708
DSNU475I   DSNUQUIB - QUIESCE UTILITY COMPLETE, ELAPSED TIME= 00:00:25
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

Chapter 20. REBUILD INDEX

REBUILD INDEX reconstructs indexes from the table that they reference.

You must modify all RECOVER INDEX jobs from a previous release of DB2 to use REBUILD INDEX instead.

For a diagram of REBUILD INDEX syntax and a description of available options, see “Syntax and options of the control statement”. For detailed guidance on running this utility, see “Instructions for running REBUILD INDEX” on page 264

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- RECOVERDB privilege for the database
- DBADM or DBCTRL authority for the database
- SYSCTRL or SYSADM authority.

To run REBUILD INDEX STATISTICS REPORT YES, the privilege set must include the SELECT privilege on the catalog tables.

Execution phases of REBUILD INDEX: The REBUILD INDEX utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup
UNLOAD	Unloading of index entries
SORT	Sorting of unloaded index entries
BUILD	Building of indexes
SORTBLD	If you specify the SORTKEYS keyword to invoke parallel index build processing for a simple or segmented table space, or a single partition of a partitioned table space, all activities that normally occur in both the SORT and BUILD phases occur in the SORTBLD phase instead.
UTILTERM	Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

© Copyright IBM Corp. 1983, 2001

REBUILD INDEX

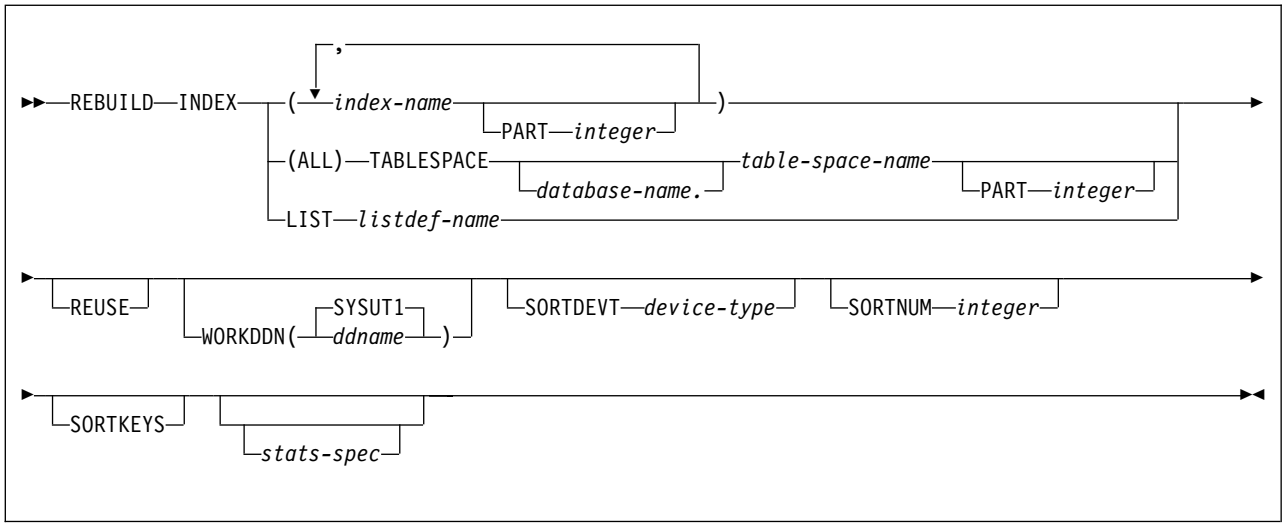

stats-spec:

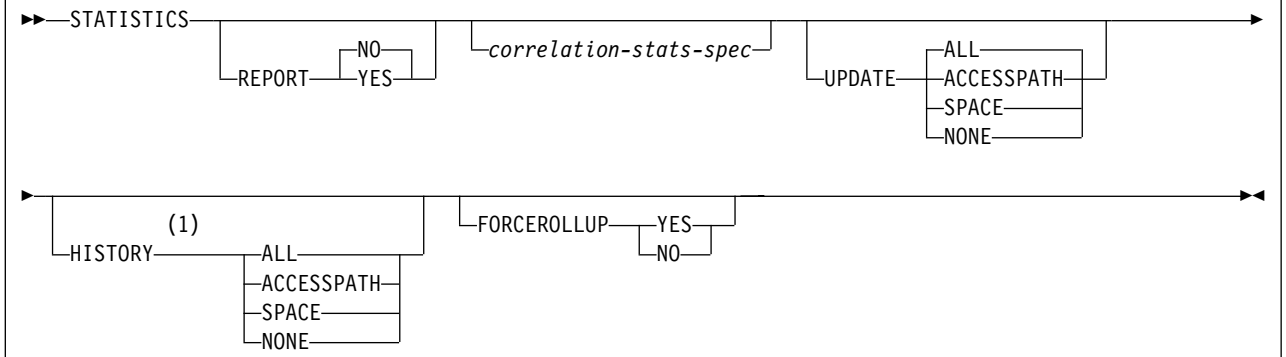

Notes:

1 What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

correlation-stats-spec:

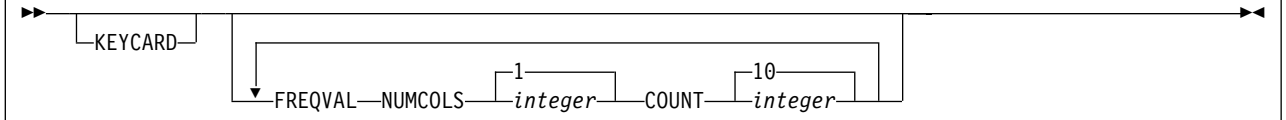

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

(*index-name***, ...)**

Indicates the qualified name of an index, in the form *creator-id.index-name*. If you omit the qualifier creator ID, the user identifier for the utility job is used.

index-name identifies the index to be rebuilt. To rebuild multiple indexes, separate each index name with a comma. All indexes listed must reside in the same table space. If more than one index is listed and TABLESPACE keyword is not specified, DB2 locates the first valid index name cited and determines the table space in which that index resides. That table space is used as the target table space for all other valid index names listed.

(ALL)
Specifies that all indexes in the table space referred to by the TABLESPACE keyword are to be rebuilt.

TABLESPACE *database-name.table-space-name*
Specifies the table space from which all indexes are to be rebuilt.

database-name
Identifies the database to which the table space belongs.

The **default** is **DSNDB04**.

table-space-name
Identifies the table space from which all indexes are rebuilt.

PART *integer*
Specifies the physical partition of a partitioning index or the logical partition of a nonpartitioning index in a partitioned table space that is to be rebuilt.

integer is the number of the partition and must be in the range from 1 to the number of partitions defined for the table space. The maximum is 254.

PART may not be specified with LIST. Use LISTDEF PARTLEVEL instead.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each control statement of the REBUILD INDEX. The list must contain either all index spaces or all table spaces. For a table space list, utility list manager will invoke REBUILD once per table space. For an index space list the utility list manager will group indexes by their related table space and execute rebuild once per table space. For more information about LISTDEF specifications, see “Generating object lists using LISTDEF” on page 144.

REUSE
Specifies that REBUILD should logically reset and reuse DB2-managed data sets without deleting and redefining them. If you do not specify REUSE, DB2 deletes and redefines DB2-managed data sets to reset them.

If you are rebuilding the index because of a media failure, do not specify REUSE.

If a data set has multiple extents, the extents will not be released if you use the REUSE parameter.

WORKDDN *ddname*
Specifies the DD statement for the temporary work file.

ddname is the DD name for the optional temporary work file.

The **default** is **WORKDDN SYSUT1**. If WORKDDN is omitted and a DD card for SYSUT1 is not provided, REBUILD INDEX performance will improve by eliminating I/O for SORT.

The WORKDDN keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a

TEMPLATE name, the utility uses DDNAME. For more information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

SORTDEVT *device-type*
Specifies the device type for temporary data sets to be dynamically allocated by DFSORT. It can be any device type acceptable to the DYNALLOC parameter of the SORT or OPTION options for DFSORT, as described in *DFSORT Application Programming: Guide*.

device-type is the device type.

A TEMPLATE specification does not dynamically allocate sort work data sets. The SORTDEVT keyword controls dynamic allocation of these data sets.

SORTNUM *integer*
Specifies the number of temporary data sets to be dynamically allocated by the sort program. If you omit SORTDEVT, SORTNUM is ignored. If you use SORTDEVT and omit SORTNUM, no value is passed to DFSORT. It is allowed to take its own default.

integer is the number of temporary data sets.

SORTKEYS
Specifies that index keys are to be sorted and built in parallel during the SORTBLD phase to improve performance, unless constrained by available memory, sort work files, or UTPRIN*nn* file allocations. If you specify SORTKEYS, utility processing ignores any WORKDDN specification or file allocation to SYSUT1 (the default WORKDDN).

STATISTICS
Specifies the gathering of index statistics.

If you specify the STATISTICS and UPDATE options, statistics are stored in the DB2 catalog. You cannot collect inline statistics for indexes on the catalog and directory tables.

REPORT
Determines if a set of messages is generated to report the collected statistics.

NO
Indicates that the set of messages is not output to SYSPRINT.

The **default** is **NO**.

YES
Indicates that the set of messages is output to SYSPRINT. The messages generated are dependent on the combination of keywords (such as TABLESPACE, INDEX, TABLE, and COLUMN) specified with the RUNSTATS utility. However, these messages are *not* dependent on the specification of the UPDATE option. REPORT YES always generates a report of SPACE and ACCESSPATH statistics.

UPDATE
Tells whether the collected statistics are inserted into the catalog tables. UPDATE also allows you to select statistics used for access path selection or statistics used by database administrators. What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

ALL Indicates that all collected statistics will be updated in the catalog.

The **default** is **ALL**.

ACCESSPATH
Indicates that only the catalog table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only the catalog table columns that provide statistics to help the database administrator assess the status of a particular table space or index are updated.

NONE Indicates that no catalog tables are updated with the collected statistics. This option is only valid when REPORT YES is specified.

HISTORY
Records all catalog table inserts or updates to the catalog history tables. What you specify under the UPDATE option affects what options you can specify under the HISTORY option. The table below lists the HISTORY options allowed for each UPDATE option specified.

The default will be supplied by the value specified in STATISTICS HISTORY on panel DSNTIPO.

ALL Indicates that all collected statistics will be updated in the catalog history tables.

ACCESSPATH
Indicates that only the catalog history table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only SPACE-related catalog statistics are updated in catalog history tables.

NONE Indicates that no catalog history tables are updated with the collected statistics.

Table 48. Allowable HISTORY options based on selected UPDATE option

UPDATE option	HISTORY options
ALL	ALL, ACCESSPATH, SPACE, NONE
ACCESSPATH	ACCESSPATH, NONE
SPACE	SPACE, NONE
NONE	NONE

FORCEROLLUP
Determines if aggregation or rollup of statistics will take place when RUNSTATS is executed in spite of some parts being empty. This will enable the optimizer to select the best access path.

The following are the available options for the FORCEROLLUP Keyword.

YES Indicates that forced aggregation or rollup processing is to be done, even though some parts may not contain data.

NO Indicates that aggregation or rollup will be done only if data is available for all parts.

If data is not available for all parts, DSNU623I message will be issued, provided the installation value for STATISTICS ROLLUP on panel DSNTIPO is set to NO.

KEYCARD
Collects all of the distinct values in all of the 1 to *n* key column combinations for the specified indexes. *n* is the number of columns in the index.

FREQVAL
Controls the collection of frequent value statistics. If you specify FREQVAL, it must be followed by two additional keywords:

NUMCOLS
Indicates the number of key columns to concatenate together when collecting frequent values from the specified index. Specifying '3' means to collect frequent values on the concatenation of the first three key columns. The **default** is **1**, which means collect frequent values on the first key column of the index.

COUNT
Indicates the number of frequent values to be collected. Specifying '15' means collect 15 frequent values from the specified key columns. The **default** is **10**.

Instructions for running REBUILD INDEX

To run REBUILD INDEX, you must:

1. Read "Before running REBUILD INDEX" in this chapter.
2. Prepare the necessary data sets, as described in "Data sets used by REBUILD INDEX".
3. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for REBUILD INDEX, see "Sample control statements" on page 272.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 265. (For a complete description of the syntax and options for REBUILD INDEX, see "Syntax and options of the control statement" on page 259.)
5. Check the compatibility table in "Concurrency and compatibility" on page 271 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the REBUILD INDEX job doesn't complete, as described in "Terminating or restarting REBUILD INDEX" on page 270.
7. Run REBUILD INDEX.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Before running REBUILD INDEX

Because the data needed to build an index is in the table space on which the index is based, you do not need image copies of indexes. To rebuild the index, you do not need to recover the table space, unless it also is damaged. Neither do you have to rebuild an index merely because you have recovered the table space it is based on.

If you recover a table space to a prior point in time and do not recover all the indexes to the same point in time, you must rebuild all of the indexes.

Data sets used by REBUILD INDEX

Table 49 on page 265 describes the data sets used by REBUILD INDEX. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 49. Data sets used by REBUILD INDEX

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
Work data sets	Temporary data sets for sort input and output. The DD names have the form SORTWK*nn*.	Yes
Work data set	Temporary data set used to store the index keys for REBUILD INDEX. The WORKDDN option of the utility control statement specifies its DD name. The default DD name is SYSUT1. To find the approximate size in bytes of the work data set, see page 265.	No

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table space
Object to be rebuilt. It is named in the REBUILD INDEX control statement and is accessed through the DB2 catalog.

Creating the work data set: REBUILD INDEX can use a single sequential data set as described by the DD statement specified in the WORKDDN option.

To calculate the approximate size (in bytes) of the WORKDDN data set, follow these steps:

1. For each table, multiply the number of records in the table by the number of indexes needing to be rebuilt on the table.
2. Add the products obtained in step 1.
3. Multiply the sum (from step 2) by the longest key length plus 9.

Allocating twice the space used by the input data sets is usually adequate for the sort work data sets. Two or three large SORTWK*nn* data sets are preferable to several small ones.

Creating the control statement

See “Syntax and options of the control statement” on page 259 for syntax and option descriptions. See “Sample control statements” on page 272 for examples of usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

“Rebuilding index partitions” on page 266
“Improving performance and space utilization” on page 266
“Building indexes in parallel for REBUILD INDEX” on page 266
“Resetting the REBUILD-pending status” on page 269
“Rebuilding critical catalog indexes” on page 270
“Rebuilt index recoverability” on page 270

Rebuilding index partitions

REBUILD INDEX can rebuild one or more partitions of a partitioning index. The PART option allows you to specify a particular partition to be rebuilt. This prevents REBUILD INDEX from unnecessarily scanning the entire table space, and unnecessarily rebuilding every index. If you recover any part of a partitioned table space to the latest point of consistency, you must rebuild all nonpartitioning indexes.

Improving performance and space utilization

REBUILD INDEX rebuilds indexes by re-creating them from the tables on which they are based. It can rebuild one or more partitions of a partitioning index in a partitioned table space when you specify the PART option. This prevents REBUILD INDEX from unnecessarily scanning the entire table space when rebuilding only a single partition of an index.

To rebuild several indexes simultaneously and reduce recovery time, use parallel index rebuild, or submit multiple index jobs. See “Building indexes in parallel for REBUILD INDEX” for more information.

When rebuilding nonpartitioning indexes and partitions of partitioning indexes, this type of parallel processing on the same table space decreases the size of the sort data set, as well as the total time required to sort all the keys.

When you run the REBUILD INDEX utility concurrently on separate partitions of a partitioning index, the sum of the processor time will be roughly equivalent to the time it takes to run a single REBUILD INDEX job against the entire index. For partitioning indexes, the elapsed time for running concurrent REBUILD INDEX jobs will be a fraction of the elapsed time for running a single REBUILD INDEX job against an entire index.

REBUILD INDEX utility performance can be improved by eliminating the work data set; however, if the job terminates abnormally, you will have to restart it from the beginning.

Building indexes in parallel for REBUILD INDEX

Use parallel index build to reduce the elapsed time for a REBUILD INDEX job by sorting the index keys and rebuilding multiple indexes in parallel, rather than sequentially. Optimally, a pair of subtasks processes each index; one subtask sorts extracted keys, while the other subtask builds the index. REBUILD INDEX begins building each index as soon as the corresponding sort generates its first sorted record. If you specified STATISTICS, a third subtask collects the sorted keys and updates the catalog table in parallel.

Note: The subtasks used for the parallel REBUILD INDEX processing use DB2 connections. If you receive message DSNU397I indicating the REBUILD INDEX utility is constrained, increase the number of concurrent connections using the MAX BATCH CONNECT parameter on panel DSNTIPE.

The greatest elapsed processing time improvements result from parallel rebuilding for:

- Multiple indexes on a simple, segmented, or partitioned table space
- The partitioning index on all partitions of a partitioned table space
- The nonpartitioning index on a partitioned table space

Figure 17 on page 267 shows a REBUILD INDEX flow with parallel index build, which requires SORTKEYS. DB2 starts multiple subtasks to unload all parts of the partitioned table space. Subtasks then sort index keys and build the partitioning

index in parallel. If you specified STATISTICS, additional subtasks collect the sorted keys and update the catalog table in parallel, eliminating the need for a second scan of the index by a separate RUNSTATS job.

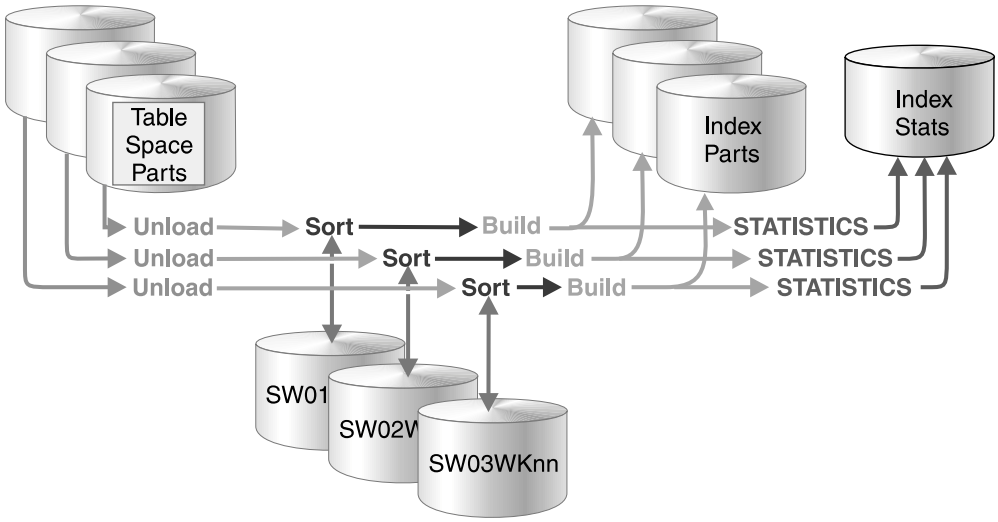

Figure 17. Rebuilding a partitioning index using parallel index build

Figure 18 shows a REBUILD INDEX flow with parallel index build, which requires SORTKEYS. DB2 starts multiple subtasks to unload all partitions of a partitioned table space, and sort index keys in parallel. The keys are then merged and passed to the build subtask, which builds the nonpartitioning index. If you specified STATISTICS, a separate subtask collects the sorted keys and updates the catalog table.

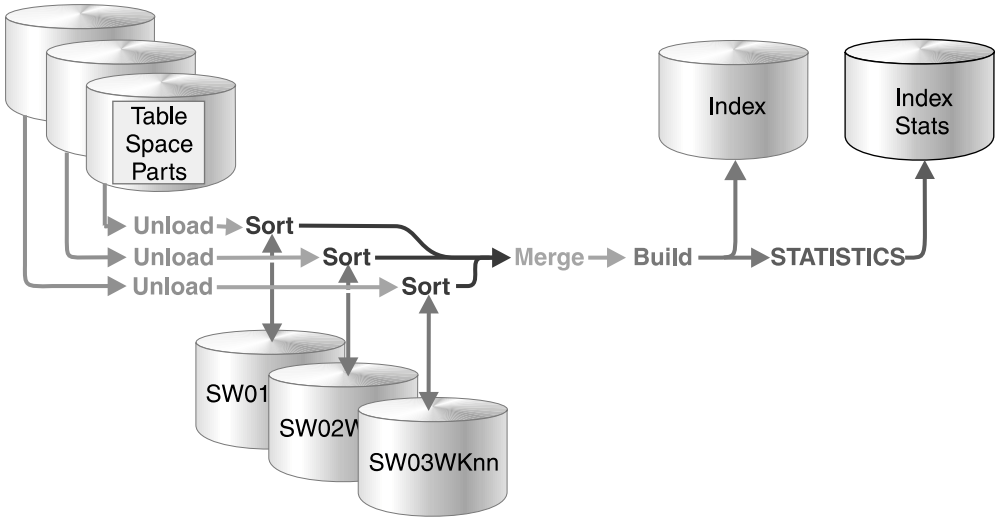

Figure 18. Rebuilding a nonpartitioning index using parallel index build

REBUILD INDEX of a table space or partition uses parallel index build if all of the following conditions are true:

- You specify the SORTKEYS keyword in the utility statement.
- You either allow the utility to dynamically allocate the data sets needed by SORT, or provide the necessary data sets yourself.

Select one of the following methods to allocate sort work and message data sets:

Method 1: REBUILD INDEX determines the optimal number of sort work and message data sets.

1. Specify the SORTKEYS and SORTDEVT keywords in the utility statement.
2. Allow dynamic allocation of sort work data sets by *not* supplying SORTWK*nn* DD statements in the REBUILD INDEX utility JCL.
3. Allocate UTPRINT to SYSOUT.

Method 2: You control allocation of sort work data sets, and REBUILD INDEX allocates message data sets.

1. Specify the SORTKEYS keyword in the utility statement.
2. Provide DD statements with DDNAMEs in the form SW*nn*WK*mm*.
3. Allocate UTPRINT to SYSOUT.

Method 3: You have the most control over rebuild processing; you must specify both sort work and message data sets.

1. Specify the SORTKEYS keyword in the utility statement.
2. Provide DD statements with DDNAMEs in the form SW*nn*WK*mm*.
3. Provide DD statements with DDNAMEs in the form UTPRIN*nn*.

Data sets used: If you select Method 2 or 3 above, define the necessary data sets by using the information provided here, along with:

- "Determining the number of sort subtasks" on page 269
- "Allocation of sort subtasks" on page 269
- "Estimating the sort work file size" on page 269

Each sort subtask must have its own group of sort work data sets and its own print message data set. In addition, you need to allocate the merge message data set when you build a single nonpartitioning index on a partitioned table space.

Possible reasons to allocate data sets in the utility job JCL rather than using dynamic allocation are to:

- Control the size and placement of the data sets
- Minimize device contention
- Optimally utilize DASD free space
- Limit the number of utility subtasks used to build indexes

The DDNAMEs SW*nn*WK*mm* define the sort work data sets that are used during utility processing. *nn* identifies the subtask pair, and *mm* identifies one or more data sets that are to be used by that subtask pair. For example:

SW01WK01 The first sort work data set that is used by the subtask building the first index.

SW01WK02 The second sort work data set that is used by the subtask building the first index.

SW02WK01 The first sort work data set that is used by the subtask building the second index.

SW02WK02 The second sort work data set that is used by the subtask building the second index.

The DDNAMEs UTPRIN*nn* define the sort work message data sets that are used by the utility subtask pairs. *nn* identifies the subtask pair.

If you allocate the UTPRINT DD statement to SYSOUT in the job statement, the sort message data sets and the merge message data set, if required, are dynamically allocated. If you want the sort, merge message data sets, or both allocated to a disk or tape data set rather than to SYSOUT, you must supply the UTPRIN*nn* or the UTMERG01 DD statements (or both) in the utility JCL. If you do not allocate the UTPRINT DD statement to SYSOUT, and do not supply UTMERG01 DD statement in the job statement, partitions are not unloaded in parallel.

Determining the number of sort subtasks: The maximum number of utility subtasks started for parallel index build is equal to the number of indexes that are to be built for a simple table space, segmented table space, or single partition of a partitioned table space, or the number of partitions that are to be unloaded if only a single index is being built on a partitioned table space.

REBUILD INDEX determines the number of subtasks according to the following guidelines:

- The number of subtasks equals the number of allocated sort work data set groups.
- The number of subtasks equals the number of allocated message data sets.
- If you allocate both sort work and message data set groups, the number of subtasks equals the smallest number of allocated data sets.

Allocation of sort subtasks: REBUILD INDEX attempts to assign one sort subtask for each index that is to be built. If REBUILD INDEX cannot start enough subtasks to build one index per subtask, it allocates any excess indexes across the pairs (in the order that the indexes were created), so that one or more subtasks might build more than one index.

Estimating the sort work file size: If you choose to provide the data sets, you need to know the size and number of keys that are present in all of the indexes being processed by the subtask in order to calculate each sort work file size. When you've determined which indexes are assigned to which subtask pairs, use the following formula to calculate the space required:

2 × (longest index key + 9) × (number of keys extracted)

longest key	The length of the longest index key that is to be processed by the subtask. For the first subtask pair for REBUILD INDEX, use the maximum value of the longest key.
number of keys	The number of keys from all indexes to be sorted that are to be processed by the subtask.

Resetting the REBUILD-pending status

REBUILD-pending (RBDP in DISPLAY command output) means that the physical or logical index partition, nonpartitioning index, or logical partition of a nonpartitioning index is in REBUILD-pending status.

There are three variations of REBUILD-pending:

RBDP The physical or logical index partition is in the REBUILD-pending status. The individual physical or logical index partition is inaccessible. RBDP is reset by rebuilding the single affected partition.

RBDP*
The logical partition of the nonpartitioning index is in the REBUILD pending status. The entire nonpartitioning index is inaccessible. RBDP* is reset by rebuilding only the affected logical partitions.

PSRBD
The nonpartitioning index space is in the REBUILD-pending status. The entire index space is inaccessible and must be rebuilt with the REBUILD INDEX utility.

You can reset the REBUILD-pending status for an index with any of these operations:
- REBUILD INDEX
- REORG INDEX SORTDATA
- REPAIR SET INDEX with NORBDPEND
- -START DATABASE with ACCESS FORCE

Important
ACCESS FORCE should be used only as a means of last resort.

Rebuilding critical catalog indexes
An ID with a granted authority receives message DSNT500I, “RESOURCE UNAVAILABLE”, while trying to rebuild indexes in the catalog or directory if the DSNDB06.SYSDBASE or DSNDB06.SYSUSER table space is unavailable. If you get this message, you must either make these table spaces available or run the RECOVER TABLESPACE utility on the catalog or directory using an authorization ID with the installation SYSADM or installation SYSOPR authority.

Rebuilt index recoverability
When you successfully rebuild an index that was defined with COPY YES, utility processing inserts a SYSIBM.SYSCOPY row with ICTYPE=’B’ for each index rebuilt. Rebuilt indexes are also placed in informational COPY-pending status. We recommend taking a full image copy of the index to create a recoverable point in time; this action also resets the ICOPY status.

Terminating or restarting REBUILD INDEX
You can terminate REBUILD INDEX with the TERM UTILITY command. If you
terminate a REBUILD INDEX job, the index space is placed in the
REBUILD-pending status and is unavailable until it has been successfully rebuilt.

If you specified the WORKDDN keyword, you can restart REBUILD INDEX during the UNLOAD and SORT phases, and at the last index built during the BUILD phase. However, there is a short period of time during writing of SORT output at the end of the SORT phase that requires restart to begin at the beginning of the UNLOAD phase instead of at the SORT phase. If you omit WORKDDN, the job starts over again from the beginning.

If you restart a job that used the SORTKEYS keyword, you must restart from the beginning of the UNLOAD phase.

If you restart a job which uses the STATISTICS keyword, inline statistics collection will not occur. To update catalog statistics, run the RUNSTATS utility after the restarted REBUILD INDEX job completes.

For more guidance in restarting online utilities, see “Restarting an online utility” on page 46 .

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Table 50 shows which claim classes REBUILD INDEX drains and any restrictive state the utility sets on the target object.

Table 51 shows which utilities can run concurrently with REBUILD INDEX on the same target object. The target object can be an index space or a partition of an index space. If compatibility depends on particular options of a utility, that is also shown.

Table 50. Claim classes of REBUILD INDEX operations. Use of claims and drains; restrictive states set on the target object.

Target	REBUILD INDEX	REBUILD INDEX PART
Table space or partition	DW/UTRO	DW/UTRO
Index or physical partition	DA/UTUT	DA/UTUT
Nonpartitioning index	DA/UTUT	DR
Logical partition of an index		DA/UTUT

Legend:
- DA - Drain all claim classes - no concurrent SQL access
- DW - Drain the write claim class - concurrent access for SQL readers
- DR - Drains the “RR” claim class
- UTUT - Utility restrictive state - exclusive control
- UTRO - Utility restrictive state - read only access allowed.

REBUILD INDEX does not set a utility restrictive state if the target object is DSNDB01.SYSUTILX.

Table 51. REBUILD INDEX compatibility

Action	REBUILD INDEX
CHECK DATA	No
CHECK INDEX	No
CHECK LOB	Yes
COPY INDEX	No
COPY TABLESPACE SHRLEVEL REFERENCE	Yes
COPY TABLESPACE SHRLEVEL CHANGE	No
DIAGNOSE	Yes
LOAD	No
MERGECOPY	Yes
MODIFY	Yes
QUIESCE	No

Table 51. REBUILD INDEX compatibility (continued)

Action	REBUILD INDEX
REBUILD INDEX	No
RECOVER INDEX	No
RECOVER TABLESPACE	No
REORG INDEX	No
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No
REORG TABLESPACE UNLOAD ONLY or EXTERNAL without cluster index	Yes
REORG TABLESPACE UNLOAD ONLY or EXTERNAL with cluster index	No
REPAIR LOCATE by KEY	No
REPAIR LOCATE by RID DUMP or VERIFY	Yes
REPAIR LOCATE by RID DELETE or REPLACE	No
REPAIR LOCATE TABLESPACE PAGE DUMP or VERIFY	Yes
REPAIR LOCATE INDEX PAGE DUMP or VERIFY	No
REPAIR LOCATE TABLESPACE or INDEX PAGE REPLACE	No
REPORT	Yes
RUNSTATS INDEX	No
RUNSTATS TABLESPACE	Yes
STOSPACE	Yes
UNLOAD	Yes

To run on SYSIBM.DSNLUX01 or SYSIBM.DSNLUX02, REBUILD INDEX must be the only utility in the job step and the only utility running in the DB2 subsystem.

Sample control statements

Example 1: Rebuild an index. Rebuild the DSN8710.XDEPT1 index, which indexes the DSN8710.TDEPT table in the DSN8D71A database.

```
//STEP1    EXEC DSNUPROC,UID='IUIQU2UT.RBLD1',TIME=1440,
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//SYSREC   DD DSN=IUIQU2UT.RBLD1.STEP1.SYSREC,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(8000,(20,20),,,ROUND)
//SYSUT1   DD DSN=IUIQU2UT.RBLD1.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(8000,(20,20),,,ROUND)
//SYSIN    DD *
REBUILD INDEX (DSN8710.XDEPT1)
//*
```

Example 2: Rebuild index partitions. Rebuild partitions 2 and 3 of the DSN8710.XEMP1 index.

```
REBUILD INDEX (DSN8710.XEMP1 PART 2, DSN8710.XEMP1 PART 3)
```

Example 3: Rebuild a single index on a segmented table space. Rebuild the DSN8710.XDEPT1 index. This example specifies the SORTKEYS keyword to use parallelism and uses dynamic data set and message set allocation with the SORTDEVT and SORTNUM keywords.

DB2 starts one utility sort subtask pair to build the index. This example does not require UTPRIN*nn* DD statements, because it uses DSNUPROC to invoke utility processing, which includes a DD statement that allocates UTPRINT to SYSOUT.

```
//SAMPJOB  JOB  ...
//STEP1    EXEC DSNUPROC,UID='SAMPJOB.RCVINDEX',UTPROC='',SYSTEM='V71A''
//SYSIN    DD *
REBUILD INDEX (DSN8710.XDEPT1)
 SORTDEVT SYSWK
 SORTNUM  4
 SORTKEYS
/*
```

Example 4: Rebuild multiple partitions of a partitioning index. Rebuild partitions 2 and 3 of the DSN8710.XDEPT1 index, using parallel index build processing. This example specifies the SORTKEYS keyword to use parallelism and uses dynamic data set and message set allocation with the SORTDEVT and SORTNUM keywords.

If sufficient virtual storage resources are available, DB2 starts one utility sort subtask pair for each partition. This example does not require UTPRIN*nn* DD statements, because it uses DSNUPROC to invoke utility processing, which includes a DD statement allocating UTPRINT to SYSOUT.

```
//SAMPJOB  JOB  ...
//STEP1    EXEC DSNUPROC,UID='SAMPJOB.RBINDEX',UTPROC='',SYSTEM='V71A'
//SYSIN    DD *
REBUILD INDEX (DSN8710.XEMP1 PART 2, DSN8710.XEMP1 PART 3)
 SORTDEVT SYSWK
 SORTNUM  4
 SORTKEYS
/*
```

Example 5: Rebuild all partitions of a partitioning index. Rebuilds all index partitions of the DSN8710.XEMP1 partitioning index, using parallel index build processing. This example specifies the SORTKEYS keyword and allocates sort work data sets in two groups, which limits the number of utility subtask pairs to two. This example does not require UTPRIN*nn* DD statements, because it uses DSNUPROC to invoke utility processing, which includes a DD statement allocating UTPRINT to SYSOUT.

```
//SAMPJOB  JOB  ...
//STEP1    EXEC DSNUPROC,UID='SAMPJOB.RCVINDEX',UTPROC='',SYSTEM='DSN'
//* First group of sort work data sets for parallel index rebuild
//SW01WK01 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SW01WK02 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SW01WK03 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//* Second group of sort work data sets for parallel index rebuild
//SW02WK01 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SW02WK02 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SW02WK03 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SYSIN    DD *
  REBUILD INDEX (DSN8710.XEMP1)
   SORTKEYS
/*
```

Example 6: Rebuild all indexes of a partitioned table space. Rebuild all indexes for table space DSN8S71E in database DSN8D71A, using parallel index build

processing. This example specifies the SORTKEYS keyword and uses dynamic data set and message set allocation with the SORTDEVT and SORTNUM keywords.

If sufficient virtual storage resources are available, DB2 starts one utility sort subtask to build the partitioning index and another utility sort subtask to build the non-partitioning index. This example does not require UTPRIN*nn* DD statements, because it uses DSNUPROC to invoke utility processing, which includes a DD statement allocating UTPRINT to SYSOUT.

```
//SAMPJOB  JOB  ...
//STEP1    EXEC DSNUPROC,UID='SAMPJOB.RCVINDEX',UTPROC='',SYSTEM='DSN'
//SYSIN    DD *
REBUILD INDEX (ALL) TABLESPACE DSN8D71A.DSN8S71E
 SORTKEYS
 SORTDEVT SYSWK
 SORTNUM  4
/*
```

Chapter 21. RECOVER

The RECOVER online utility recovers data to the current state or to a previous point in time by restoring a copy, then applying log records.

The largest unit of data recovery is the table space or index space; the smallest is the page. You can recover a single object, or a list of objects. RECOVER recovers an entire table space, index space, a partition or data set, pages within an error range, or a single page. You recover data from image copies of an object and from log records containing changes to the object. If the most recent full image copy data set is unusable, and there are previous image copy data sets existing in the system, then RECOVER uses the previous image copy data sets. For more information about using data sets for recovery, see “Recovering a data set or partition” on page 284.

Compatibility with prior releases: In previous releases of DB2, REBUILD INDEX was called RECOVER INDEX. You must modify all utility control statements from previous releases to use REBUILD INDEX if you want to continue recovering the indexes via a scan of the data. However, if you want to recover the indexes from a full image copy, change those control statements to use the new RECOVER INDEX syntax on page 276. Only indexes that were defined with the COPY YES attribute can be copied and subsequently recovered; see Chapter 5 of *DB2 SQL Reference* for more information about the COPY YES attribute of the ALTER INDEX and CREATE INDEX SQL statements.

For a diagram of RECOVER syntax and a description of available options, see “Syntax and options of the control statement” on page 276. For detailed guidance on running this utility, see “Instructions for running RECOVER” on page 282.

Output: Output from RECOVER consists of recovered data (either a table space, index, partition or data set, error range, or page within a table space).

If you specify the TOLOGPOINT, TORBA, or TOCOPY option to recover data to a point in time, RECOVER puts any associated index spaces in REBUILD-pending status. You must run REBUILD INDEX to remove the index space from REBUILD-pending status.

If you use the RECOVER utility to partially recover a referentially-related table space set or a base table space and LOB table space set, you must ensure that you recover the entire set of table spaces, including rebuilding or recovering all indexes (including indexes on auxiliary tables for a base table space and LOB table space set), to a common quiesce point or a SHRLEVEL REFERENCE copy. If you do not include every member of the set, or if you do not recover the entire set to the same point in time, RECOVER sets the CHECK-pending status on for all dependent table spaces, base table spaces, or LOB table spaces in the set.

If you use the RECOVER utility to partially recover data and all indexes over the data, it is best to recover these objects to a common quiesce point or SHRLEVEL REFERENCE copy. Otherwise, RECOVER places all indexes in the CHECK-pending status.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- RECOVERDB privilege for the database
- DBADM or DBCTRL authority for the database

© Copyright IBM Corp. 1983, 2001

- SYSCTRL or SYSADM authority.

An ID with installation SYSOPR authority can also execute RECOVER, but only on a table space in the DSNDB01 or DSNDB06 database.

Execution phases of RECOVER: The RECOVER utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup
RESTORE	Locate and merge any appropriate image copies and restore the table space to a backup level; processes a list of objects in parallel if you specified the PARALLEL keyword.
RESTORER	If you specified the PARALLEL keyword, reads and merges the image copies.
RESTOREW	If you specified the PARALLEL keyword, writes the pages to the object.
LOGAPPLY	Apply any outstanding log changes to the object restored from the previous phase or step.
UTILTERM	Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

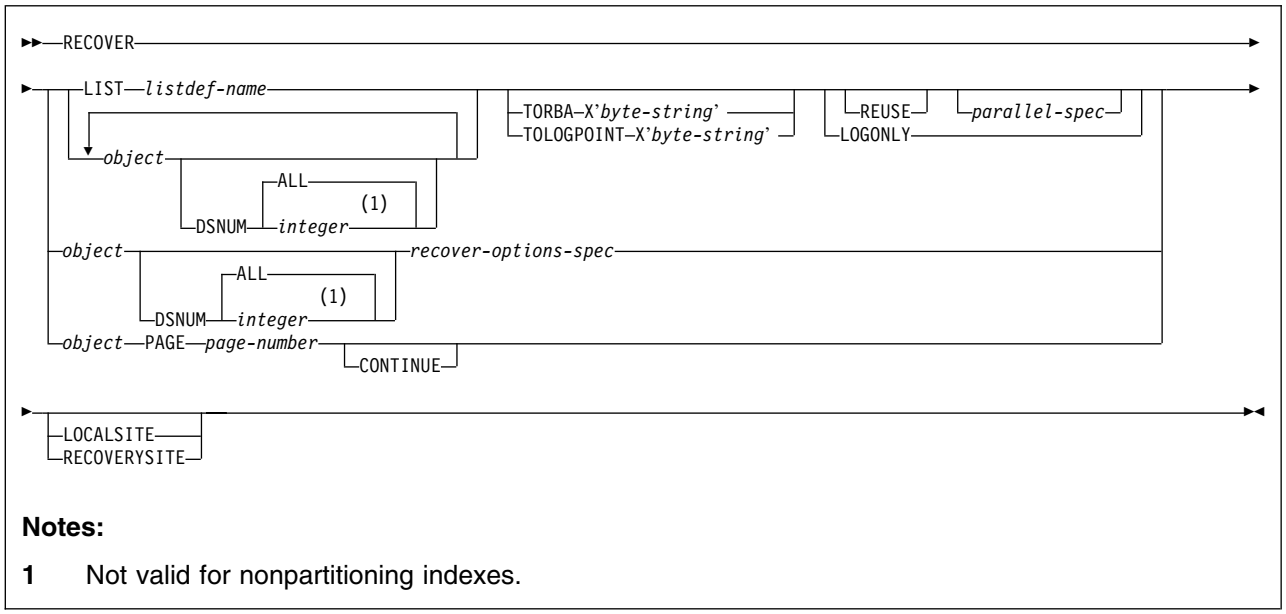

Notes:

1 Not valid for nonpartitioning indexes.

object:

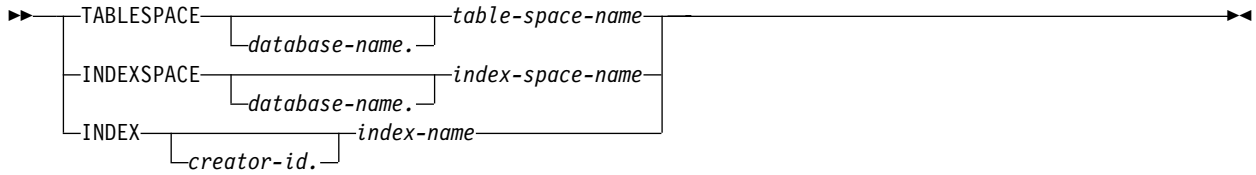

parallel-spec:

recover-options-spec:

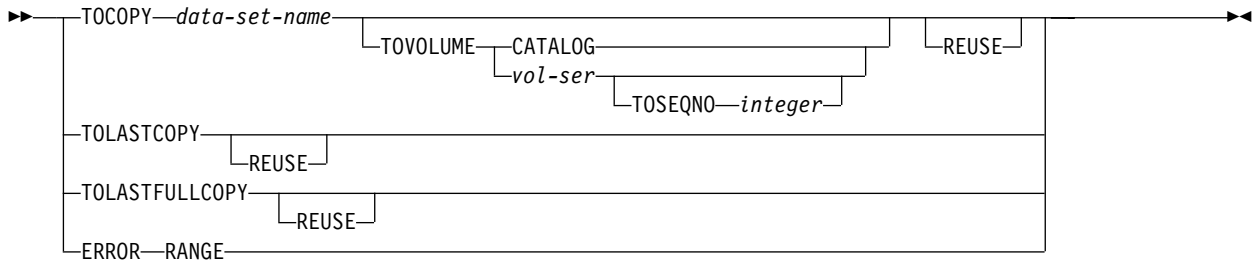

Option descriptions

You can specify a list of objects by repeating the TABLESPACE, INDEX, or INDEXSPACE keywords. If you use a list of objects, the valid keywords are: DSNUM, TORBA, TOLOGPOINT, LOGONLY, PARALLEL, and either LOCALSITE or RECOVERYSITE.

RECOVER cannot recover a table space or index space that is defined to use a storage group that is defined with mixed specific and nonspecific volume IDs. If you specify such a table space or index space, the job terminates and you receive error message DSNU419I.

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each control statement of RECOVER. The list can contain a mixture of table spaces and index spaces. Do not specify LIST with

the TABLESPACE, INDEX, INDEXSPACE, PAGE, TOCOPY, ERROR RANGE, or DSNUM keyword. Utility list manager will invoke RECOVER once for the entire list.

For more information about LISTDEF specifications, see "Generating object lists using LISTDEF" on page 144.

TABLESPACE *database-name.table-space-name*
Specifies the table space (and optionally, the database to which it belongs) that is to be recovered.

You can specify a list of table spaces by repeating the TABLESPACE keyword. You can recover an individual catalog or directory table space in a list with its IBM-defined indexes. You cannot recover multiple catalog or directory table spaces in a list.

database-name
Is the name of the database the table space belongs to.

The **default** is **DSNDB04**.

table-space-name
Is the name of the table space to be recovered.

INDEXSPACE *database-name.index-space-name*
Specifies the index space that is to be recovered.

database-name
Specifies the name of the database the index space belongs to.

The **default** is **DSNDB04**.

index-space-name
Specifies the name of the index space to be recovered.

INDEX *creator-id.index-name*
Specifies the index in the index space that is to be recovered. The RECOVER utility can recover only indexes that were defined with the COPY YES attribute and subsequently copied. For more information about this restriction, see "Compatibility with Prior Releases" on page 275.

creator-id Optionally specifies the creator of the index.

The **default** is the user identifier for the utility.

index-name Specifies the name of the index in the index space to be recovered.

The following are optional:

DSNUM
Identifies a partition or data set within a table space, or a partition within an index space, that is to be recovered; or it recovers the entire table space or index space.

ALL
Recovers the entire table space or index space.

The **default** is **ALL**.

integer
Is the number of a partition or data set to be recovered. The maximum is 254.

Not valid for non-partitioning indexes.

For a partitioned table space or index space: The integer is its partition number.

For a nonpartitioned table space: Find the integer at the end of the data set name as cataloged in the VSAM catalog. The data set name has this format:

catname.DSNDB*x*.*dbname*.*tsname*.*y*0001.A*nnn*

where:

catname	The VSAM catalog name or alias
x	C or D
dbname	The database name
tsname	The table space name
y	I or J
nnn	The data set integer.

PAGE *page-number*
Specifies a particular page to be recovered.

page-number is the number of the page, in either decimal or hexadecimal notation. For example, both 999 and X'3E7' represent the same page. PAGE is invalid with the LIST specification.

CONTINUE
Specifies that the recovery process is to continue. Use this option only if RECOVER has terminated during reconstruction of a page, because of an error. In this case, the page is marked as “broken”. After you have repaired the page, you can use the CONTINUE option to recover the page, starting from the point of failure in the recovery log.

TORBA X'*byte-string***'**
Is used in a non-data-sharing environment to specify a point on the log to recover to. You must specify TORBA when recovering to a point before Version 4. Specify an RBA value.

In a data sharing environment, TORBA must be used only when recovering to a point before the originating member joined the data sharing group. If you specify an RBA after this point, the recovery fails.

Terminates the recovery process with the last log record whose relative byte address (RBA) is not greater than *byte-string*, which is a string of up to 12 hexadecimal characters. If *byte-string* is the RBA of the first byte of a log record, that record is included in the recovery.

TOLOGPOINT X'*byte-string***'**
Specifies a point on the log to recover to. Specify either an RBA or an LRSN value.

The LRSN is a string of 12 hexadecimal characters and is reported by the DSN1LOGP utility. The value must be greater than the LRSN value when Version 4 started.

LOGONLY
Recovers the target objects from their existing data sets by applying only log records to the data sets. DB2 applies all log records that were written after a point that is recorded in the data set itself.

Use the LOGONLY option when the data sets of the target objects have already been restored to a point of consistency by another process offline, such as DFSMS Concurrent Copy.

When recovering an index space using RECOVER LOGONLY, use the COPY
YES attribute. Otherwise, RECOVER terminates with a message DSNU425I.

REUSE
Specifies that RECOVER logically resets and reuses DB2-managed data sets without deleting and redefining them. If you do not specify REUSE, DB2 deletes and redefines DB2-managed data sets to reset them.

If you are recovering an object because of a media failure, do not specify REUSE.

If a data set has multiple extents, the extents will not be released if you use the REUSE parameter.

PARALLEL
Specifies the maximum number of objects in the list that should be restored in parallel from image copies on DASD. Specify the PARALLEL keyword to take advantage of parallel processing during the RESTORE phase.

(*num-objects***)**
Specifies the number of objects in the list that should be processed in parallel. If storage constraints are encountered, you can adjust this value to a smaller value.

If you specify 0 or do not specify a value for *num-objects*, RECOVER determines the optimal number of objects to process in parallel.

LOCALSITE
RECOVER uses image copies from the local site. If you specify neither LOCALSITE or RECOVERYSITE, then RECOVER uses image copies from the current site of invocation. (The current site is identified on the installation panel DSNTIPO under SITE TYPE and in the macro DSN6SPRM under SITETYP.)

RECOVERYSITE
RECOVER uses image copies from the recovery site. If you specify neither LOCALSITE or RECOVERYSITE, then RECOVER uses image copies from the current site of invocation. (The current site is identified on the installation panel DSNTIPO under SITE TYPE and in the macro DSN6SPRM under SITETYP.)

TOCOPY *data-set-name*
Specifies the particular image copy data set that DB2 uses as a source for recovery.

data-set-name is the name of the data set.

If the data set is a full image copy, it is the only data set used in recovery. If it is an incremental image copy, the recovery also uses the previous full image copy and any intervening incremental image copies.

If you specify the data set as the local backup copy, DB2 first tries to allocate the local primary copy. If the local primary copy is unavailable, DB2 uses the local backup copy.

If you use TOCOPY or TORBA to recover a single data set of a nonpartitioned table space, DB2 issues message DSNU520I to warn that the table space can become inconsistent following the RECOVER job. This point in time recovery can cause compressed data to exist without a dictionary or can even overwrite the data set that contains the current dictionary.

If you use TOCOPY with a particular partition or data set (identified with DSNUM), then the image copy must be for the same partition or data set, or for the whole table space or index space. If you use TOCOPY with DSNUM ALL, the image copy must be for DSNUM ALL.

If the image copy data set is an MVS generation data set, then supply a fully qualified data set name including the absolute generation and version number.

TOCOPY is invalid with a LIST specification.

If the image copy data set is not a generation data set and there is more than one image copy data set with the same data set name, use one of the following options to identify the data set exactly:

TOVOLUME
Identifies the image copy data set.

CATALOG
Identifies the data set as cataloged. Use this option only for an image copy that was created as a cataloged data set. (Its volume serial is not recorded in SYSIBM.SYSCOPY.)

RECOVER refers to the SYSIBM.SYSCOPY catalog table during execution. If you use TOVOLUME CATALOG, the data set must be cataloged. If you remove the data set from the catalog after creating it, you must catalog the data set again to make it consistent with the record for this copy that appears in SYSIBM.SYSCOPY.

vol-ser
Identifies the data set by an alphanumeric volume serial identifier of its first volume. Use this option only for an image copy that was created as a noncataloged data set. Specify the first *vol-ser* in the SYSCOPY record to locate a data set stored on multiple tape volumes.

TOSEQNO *integer*
Identifies the image copy data set by its file sequence number.

integer is the file sequence number.

TOLASTCOPY
Specifies that the Recover Utility should restore the object to the last copy taken. If the last copy taken is a full image copy, then it is restored to the object. If the last copy taken is an incremental image copy, then the most recent full copy along with any incremental copies are restored to the object.

TOLASTFULLCOPY
Specifies that the Recover Utility should restore the object to the last full image copy that was taken. Any incremental image copies that were taken after the full copy are not restored to the object.

ERROR RANGE
Specifies that all pages within the range of reported I/O errors are to be recovered. Recovering an error range is useful when the range is small relative to the object containing it; otherwise, it is probably better to recover the entire object.

There are some situations in which recovery using the ERROR RANGE option is not possible, such as when a sufficient quantity of alternate tracks cannot be obtained for all bad records within the error range. The IBM Device Support Facility, ICKDSF service utility, can be used to determine whether this situation exists. In such a situation, the error data set should be redefined at a different location on the volume or on a different volume and the RECOVER utility will run without the ERROR RANGE option.

ERROR RANGE is invalid with a LIST specification.

Refer to Part 4 (Volume 1) of *DB2 Administration Guide* for additional information concerning the use of this keyword.

Instructions for running RECOVER

To run RECOVER, you must:

1. Read “Before running RECOVER” in this chapter.
2. Prepare the necessary data sets, as described in “Data sets used by RECOVER” on page 283.
3. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for RECOVER, see “Sample control statements” on page 301.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 283. (For a complete description of the syntax and options for RECOVER, see “Syntax and options of the control statement” on page 276.)
5. Check the compatibility table in “Concurrency and compatibility” on page 300 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the RECOVER utility job does not complete, as described in “Terminating or restarting RECOVER” on page 299.
7. Run RECOVER.

See “Chapter 3. Invoking DB2 online utilities” on page 23 for an explanation of ways to execute DB2 utilities.

Before running RECOVER

Recovering data and indexes: You do not always need to recover both the data and indexes. If you recover the table space or index space to a current RBA or LRSN, then any referentially-related objects do not need to be recovered. If you plan to recover a damaged object to a point in time, ensure that you use a consistent point in time for all of its referentially-related objects, including related LOB table spaces. You must rebuild the indexes from the data if one of the following conditions is true:

- The table space is recovered to a point-in-time
- An index is damaged
- An index is in REBUILD-pending status
- No image copy of the index is available

If you need to recover both the data and the indexes, and no image copies of the indexes are available, use the following procedure:

1. Use RECOVER TABLESPACE to recover the data.
2. Run REBUILD INDEX on any related indexes to rebuild them from the data.

If you have image copies of both the table spaces and the indexes, you can recover both sets of objects in the same RECOVER utility statement. The objects are recovered from the image copies and logs.

If the table space or index space to be recovered is associated with a storage group, DB2 deletes and defines the necessary data sets. If the STOGROUP has been altered to remove the volume on which the table space or index space is located, RECOVER places the data set on another volume of the storage group.

| Some logging can occur if the index is a type 2 non-partitioning index and is being
| concurrently accessed either by SQL being done on a different partition of the same
| table space or by a utility being run on a different partition of the same table space.

Data sets used by RECOVER

Table 52 describes the data sets used by RECOVER. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 52. Data sets used by RECOVER

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes

The following objects are named in the utility control statement and do not require DD cards in the JCL:

Table space or Index space The name of the table space or index space to be recovered. It is named in the control statement and is accessed through the DB2 catalog. If you want to recover less than an entire table space:
- Use the DSNUM option to recover a partition or data set.
- Use the PAGE option to recover a single page.
- Use the ERROR RANGE option to recover a range of pages with I/O errors.

Image copy data set This information is accessed through the DB2 catalog. However, if you would like to preallocate your image copy data sets by using DD cards, refer to “Retaining tape mounts” on page 298 for more information.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

“Recovering a table space”
“Recovering a list of objects” on page 284
“Recovering a data set or partition” on page 284
“Recovering with incremental copies” on page 285
“Recovering a page” on page 285
“Recovering an error range” on page 286
“Recovering with a data set copy not made by DB2” on page 286
“Recovering catalog and directory objects” on page 287
“Recovering a table space containing LOB data” on page 290
“Performing a point-in-time recovery” on page 291
“Avoiding specific image copy data sets” on page 295
“Improving performance” on page 295
“Optimizing the LOGAPPLY phase” on page 295
“Recovering image copies in a JES3 environment” on page 297
“Resetting RECOVER-pending or REBUILD-pending status” on page 297.

Recovering a table space

The following RECOVER statement recovers table space DSN8S71D in database DSN8D71A:

```
RECOVER TABLESPACE DSN8D71A.DSN8S71D
```

You can also recover multiple table spaces by creating a list of table spaces to be recovered; repeat the TABLESPACE keyword before each table space specified.

The following statement recovers partition 2 of the partitioned table space DSN8D71A.DSN8S71E, and recovers the table space DSN8D71A.DSN8S71D to the quiesce point (RBA X'000007425468').

```
RECOVER TABLESPACE DSN8D71A.DSN8S71E DSNUM 2
        TABLESPACE DSN8D71A.DSN8S71D
        TORBA X'000007425468'
```

Each table space is unavailable for most other applications until recovery is complete. If you make image copies by table space, you can recover the entire table space, or a data set or partition from the table space. If you make image copies separately by partition or data set, you must recover the partitions or data sets by separate RECOVER operations. The following example shows the RECOVER statement for recovering four data sets in database DSN8D71A, table space DSN8S71E:

```
RECOVER TABLESPACE DSN8D71A.DSN8S71E DSNUM 1
        TABLESPACE DSN8D71A.DSN8S71E DSNUM 2
        TABLESPACE DSN8D71A.DSN8S71E DSNUM 3
        TABLESPACE DSN8D71A.DSN8S71E DSNUM 4
```

The recovery of these data sets can be scheduled in four separate jobs to run in parallel. In many cases, the four jobs can read the log data concurrently.

If a table space or data set is in the COPY-pending status, recovering it might not be possible. This status can be reset in several ways: see “Resetting the COPY-pending status” on page 209.

Recovering a list of objects

You can recover any of the following objects:

- Table space
- Table space partition
- Piece of a linear table space
- Index space
- Index space partition

When you recover to a prior point-in-time, you should recover a set of referentially related table spaces together to avoid putting any of the table spaces in CHECK-pending status. Use REPORT TABLESPACESET to get a table space listing.

RECOVER merges incremental copies serially and dynamically. As a result, recovery of a table space list with numerous incremental copies can be time-consuming and operator-intensive.

If referential integrity violations are not an issue, you can run a separate job to recover each table space.

When you specify the PARALLEL keyword, DB2 supports parallelism during the RESTORE phase for image copies on DASD devices. If RECOVER encounters a tape volume in the list, processing of remaining objects pauses until the tape object has completed, and then parallel processing resumes.

Recovering a data set or partition

You can use the RECOVER utility to recover individual partitions and data sets. The phases for data set recovery are the same as for table space recovery.

When image copies are done at the data set level, then RECOVER must be done at the data set level. To recover the whole table space, all the data sets must be

recovered individually in one or more RECOVER steps. If recovery is attempted at the table space level, the following message is received:

```
DSNU514I DSNUCASA - RECOVERY DATA DOES NOT PERMIT
                    TABLESPACE RECOVERY
```

Alternatively, if image copies are taken at the table space, index, or index space level, individual data sets can be recovered simply by coding the DSNUM parameter.

RECOVER does *not* support recovery of:

- A single data set for non-partitioning indexes
- A logical part of a nonpartitioning index

You should consider the effects of data compression on recovery. When you use the option TOCOPY, TOLOGPOINT, or TORBA to recover a single data set of a nonpartitioned table space, message DSNU520I is issued to warn you that the table space might be inconsistent after the RECOVER job. This point-in-time recovery might cause compressed data to exist without a dictionary or could even overwrite the data set that contains the current dictionary.

Recovering with incremental copies

The RECOVER utility merges all incremental image copies since the last full image copy, and must have all the image copies available at the same time. If there is any likelihood that the requirement will strain your system resources—for example, by demanding more tape units than are available—consider running MERGECOPY regularly to merge image copies into one copy.

Even if you do not periodically merge multiple image copies into one copy when there are not enough tape units, the utility can still perform. RECOVER dynamically allocates the full image copy and attempts to allocate dynamically all the incremental image copy data sets. If RECOVER successfully allocates every incremental copy, recovery proceeds to merge pages to table spaces and apply the log. If a point is reached where an incremental copy cannot be allocated, RECOVER notes the log RBA or LRSN of the last successfully allocated data set. Attempts to allocate incremental copies cease, and the merge proceeds using only the allocated data sets. The log is applied from the noted RBA or LRSN, and the incremental image copies that were not allocated are simply ignored.

Recovering a page

RECOVER PAGE allows you to recover data on a page that has been damaged. In some situations, you can determine (usually from an error message) which page of an object has been damaged. You can use the PAGE option to recover a single page. You can use the CONTINUE option to continue recovering a page that was damaged during the LOGAPPLY phase of a RECOVER operation.

Recovering a page using PAGE and CONTINUE: Suppose you start RECOVER for table space TSPACE1. During processing, message DSNI012I informs you of a problem that damages page number 5. RECOVER completes, but the damaged page, number 5, is in a *stopped* state and is not recovered. When RECOVER ends, message DSNU501I informs you that page 5 is damaged.

To repair the damaged page:

1. Use the DUMP option of the REPAIR utility to view the contents of the damaged page. Determine what change should have been made by the applicable log record and apply it by using the REPLACE option of REPAIR. Use the RESET option to turn off the *inconsistent data* indicator.

Attention: Be extremely careful when using the REPAIR utility to replace data. Changing data to invalid values using REPLACE can produce unpredictable results, particularly when changing page header information. Improper use of REPAIR can result in damaged data, or in some cases, system failure.

2. Resubmit the RECOVER utility job specifying TABLESPACE(TSPACE1) PAGE(*5*) CONTINUE. The RECOVER utility finishes recovering the damaged page by applying the log records remaining after the one that caused the problem.

If more than one page is damaged during RECOVER, do step 1 on page 285 and step 2 for each damaged page.

Recovering an error range

The ERROR RANGE option of RECOVER allows you to repair pages with reported I/O errors. DB2 maintains a page error range for I/O errors for each data set; pages within the range cannot be accessed. The DISPLAY DATABASE command displays the range. When recovering an error range, RECOVER:

1. Locates, allocates, and applies image copies.
2. Applies changes from the log.

The following statement recovers any current error range problems for table space TS1:

```
RECOVER TABLESPACE DB1.TS1 ERROR RANGE
```

Recovering an error range is useful when the range is small relative to the object containing it; otherwise, it is probably better to recover the entire object.

Message DSNU086I indicates that I/O errors were detected on a table space and that you need to recover it. Before you attempt to use the ERROR RANGE option of RECOVER, you should run the ICKDSF service utility to correct the DASD error. If an I/O error is detected during RECOVER processing, DB2 issues message DSNU538I to tell you which target tracks are involved. The message provides enough information to run ICKDSF correctly.

There are some situations, announced by error messages, in which recovery of an error range only is not possible. In such a situation, it is better to recover the entire object.

During the recovery of the entire table space or index space, DB2 might still encounter I/O errors that indicate DB2 is still using a bad volume. For user-defined data sets, you should use access method services to delete the data sets, and redefine them with the same name on a new volume. If you use DB2 storage groups, then you can remove the bad volume from the storage group using ALTER STOGROUP.

Recovering with a data set copy not made by DB2

You can restore a data set to a point of consistency by using a data set copy that was not made by the COPY utility. After recovery to the point of consistency, if you choose to continue and recover to the current point in time, you do not want RECOVER to begin processing by restoring the data set from a DB2 image copy. Therefore, use the LOGONLY option of RECOVER, which will cause RECOVER to skip the RESTORE phase and apply the log records only, starting from the first log record written after the data set was backed up.

Because the data sets are restored offline without DB2's involvement, RECOVER LOGONLY checks that the data set identifiers match those in the DB2 catalog. If they do not match, message DSNU548I is issued and the job terminates with return code 8.

You can use the LOGONLY option on a list of objects.

To ensure that no other transactions can access DB2 objects between the time that you restore a data set and the time that you run RECOVER LOGONLY:

1. Stop the DB2 objects being recovered by issuing the following command:
   `-STOP DATABASE(database name) SPACENAM(space-name)`
2. Restore all DB2 data sets that are being recovered.
3. Start the DB2 objects being recovered by issuing the following command:
   `-START DATABASE(database name) SPACENAM(space-name) ACCESS(UT)`
4. Run the RECOVER utility without the TORBA or TOLOGPOINT parameters and with the LOGONLY parameter to recover the DB2 data sets to the current point in time and to perform forward recovery using DB2 logs. If you want to recover the DB2 data sets to a prior point in time, run the RECOVER utility with either TORBA or TOLOGPOINT, and the LOGONLY parameters.
5. If you did not recover related indexes in the same RECOVER control statement, then rebuild all indexes on the recovered object.
6. Issue the following command to allow access to the recovered object if the recovery completes successfully:
   `-START DATABASE(database name) SPACENAM(space-name) ACCESS(RW)`

With the LOGONLY option, when recovering a single piece of a multi-piece linear page set, RECOVER opens the first piece of the page set. If the data set is migrated by]DFSMShsm™, then the data set is recalled by DFSMShsm. Without LOGONLY, no data set recall is requested.

Backing up a single piece of a multi-piece linear page set is not recommended. It can cause a data integrity problem if the backup is used to restore the data set at a later time.

Recovering catalog and directory objects

If you are recovering any subset of the objects in the following list, start with the object that appears first and continue in the order of the list. For example, if you must recover SYSLGRNX, SYSUTILX, and SYSUSER, recover first SYSUTILX, then SYSLGRNX, then SYSUSER. It is not necessary to recover all of the objects, only those that require recovery.

1. DSNDB01.SYSUTILX
2. All indexes on SYSUTILX
3. DSNDB01.DBD01
4. DSNDB06.SYSCOPY
5. All IBM defined indexes on SYSCOPY[3]
6. DSNDB01.SYSLGRNX
7. All indexes on SYSLGRNX
8. DSNDB06.SYSDBAUT
9. All IBM defined indexes on SYSDBAUT[3]

3. If there are no user defined indexes on the catalog, execute `REBUILD INDEX (ALL) TABLESPACE DSNDB06.SYSxxxx` to rebuild all the IBM defined indexes on a catalog table space. If user defined indexes are created on the catalog, the IBM defined indexes must be rebuilt individually and the user defined indexes rebuilt in a subsequent step. See Appendix D of *DB2 SQL Reference* for a list of the IBM defined indexes.

10. DSNDB06.SYSUSER
11. DSNDB06.SYSDBASE
12. All IBM defined indexes on SYSDBASE and SYSUSER[3]
13. Other catalog and directory table spaces and indexes. The remaining catalog table spaces, in database DSNDB06, are SYSGROUP, SYSGPAUT, SYSOBJ, SYSPLAN, SYSPKAGE, SYSSEQ, SYSSEQ2, SYSSTATS, SYSSTR, and SYSVIEWS. Most indexes are listed in *DB2 SQL Reference*. There are two remaining directory table spaces, DSNDB01.SCT02, which has indexes SYSIBM.DSNSCT02, and DSNDB01.SPT01, which has indexes SYSIBM.DSNSPT01 and SYSIBM.DSNSPT02.
14. All user defined indexes on the catalog.
15. System utility table spaces such as QMF.
16. If used, the communications database (CDB), the object and application registration tables, and the resource limit specification tables.
17. User table spaces.

For all catalog and directory table spaces, the IBM-defined indexes with the COPY YES attribute can be listed in the same RECOVER utility statement.

Recovery of the items on the list can be done concurrently or included in the same job step. However, some restrictions apply:

1. When you recover the following table spaces or indexes, the job step in which the RECOVER statement appears must not contain any other utility statements. No other utilities can run while the RECOVER utility is running.
   - DSNDB01.SYSUTILX
   - All indexes on SYSUTILX
   - DSNDB01.DBD01
2. When you recover the following table spaces, no other utilities can run while the RECOVER utility is running. Other utility statements may exist in the same job step.
   - DSNDB06.SYSCOPY
   - DSNDB01.SYSLGRNX
   - DSNDB06.SYSDBAUT
   - DSNDB06.SYSUSER
   - DSNDB06.SYSDBASE

Attention: If the logging environment requires adding or restoring active logs, restoring archive logs, or performing any action that affects the log inventory in the BSDS, you should recover the BSDS before catalog and directory objects. For information on recovering the BSDS, see Part 4 (Volume 1) of *DB2 Administration Guide* .

The access method services REPRO function should be used to copy active log data sets. For information on the JCL for the access method services REPRO function, see:

- *DFSMS/MVS: Access Method Services for the Integrated Catalog*
- *DFSMS/MVS: Access Method Services for VSAM Catalogs*

Why the order is important: To recover one object, RECOVER must obtain information about it from some other object. Table 53 on page 289 lists the objects from which RECOVER must obtain information.

Table 53. Objects the RECOVER utility accesses

Object Name	Reason for Access by RECOVER
DSNDB01.SYSUTILX	Utility restart information. It is not accessed when it is recovered; RECOVER for this object is not restartable, and there can be no other commands in the same job step. SYSCOPY information for SYSUTILX is obtained from the log.
DSNDB01.DBD01	Descriptors for the catalog database (DSNDB06), the work file database (DSNDB07), and user databases. RECOVER for this object is not restartable, and there can be no other commands in the same job step. SYSCOPY information for DBD01 is obtained from the log.
DSNDB06.SYSCOPY	Locations of image copy data sets. SYSCOPY information for SYSCOPY itself is obtained from the log.
DSNDB01.SYSLGRNX	The RBA or LRSN of the first log record after the most recent copy.
DSNDB06.SYSDBAUT, DSNDB06.SYSUSER	To verify that the authorization ID is authorized to run RECOVER.
DSNDB06.SYSDBASE	Descriptors of table spaces to be recovered.

You can use REPORT RECOVERY to get SYSCOPY information for DSNDB01.SYSUTILX, DSNDB01.DBD01, and DSNDB06.SYSCOPY.

Planning for point-in-time recovery for the catalog and directory: When you recover the DB2 catalog and directory, consider the entire catalog and directory including all table spaces and index spaces as one logical unit. Recover all objects in the catalog and directory to the same point of consistency. If a point-in-time recovery of the catalog and directory objects is planned, a separate quiesce of the DSNDB06.SYSCOPY table space is required after a quiesce of the other catalog and directory table spaces. Recommendation: Recover the catalog and directory objects to the current state. Sample queries and documentation are provided in DSNTESQ in the SDSNSAMP sample library that can be used to check the consistency of the catalog.

Indexes are rebuilt by REBUILD INDEX. If the only items you have recovered are table spaces in the catalog or directory, you might have to rebuild their indexes. Use the CHECK INDEX utility to determine whether an index is inconsistent with the data it indexes. You can use the RECOVER utility to recover catalog and directory indexes if the index was defined with the COPY YES attribute and you have a full index image copy.

You must recover the catalog and directory before recovering user table spaces.

Be aware that the following table spaces, along with the indexes over them, do not have entries in SYSIBM.SYSLGRNX, even if they were defined with COPY YES:

- DSNDB01.SYSUTILX
- DSNDB01.DBD01
- DSNDB01.SYSLGRNX
- DSNDB06.SYSCOPY
- DSNDB06.SYSGROUP

- DSNDB01.SCT02
- DSNDB01.SPT01

These objects are assumed to be open from the point of their last image copy, so the RECOVER utility processes the log from that point forward.

Point-in-time recovery: Full recovery of the catalog and directory table spaces and indexes is strongly recommended. However, if you need to plan for point-in-time recovery of the catalog and directory, then one method of creating a point of consistency is to:

1. Quiesce all catalog and directory table spaces in a list, except for DSNDB06.SYSCOPY and DSNDB01.SYSUTILX.
2. Quiesce DSNDB06.SYSCOPY.

   We recommend that you quiesce DSNDB06.SYSCOPY in a separate utility statement; when you recover DSNDB06.SYSCOPY to its own quiesce point, it contains the ICTYPE = 'Q' (quiesce) SYSIBM.SYSCOPY records for the other catalog and directory table spaces.
3. Quiesce DSNDB01.SYSUTILX in a separate job step.

If you need to recover to a point in time, recover DSNDB06.SYSCOPY and DSNDB01.SYSUTILX to their own quiesce points, and recover other catalog and directory table spaces to their common quiesce point. The catalog and directory objects must be recovered in a particular order, as described on page 288.

Recovering critical catalog table spaces: An ID with a granted authority receives message DSNT500I, “RESOURCE UNAVAILABLE”, while trying to recover a table space in the catalog or directory if table space DSNDB06.SYSDBASE or DSNDB06.SYSUSER is unavailable. If you get this message, you must either make these table spaces available or run the RECOVER utility on the catalog or directory using an authorization ID with the installation SYSADM or installation SYSOPR authority.

Recovering a table space containing LOB data

The RECOVER utility can set the auxiliary warning status for a LOB table space if it finds at least one invalid LOB column. A LOB can be marked invalid if all of the following are true:

1. The LOB table space was defined with the LOG(NO) attribute.
2. The LOB table space was recovered.
3. The LOB was updated since the last image copy.

The status of an object related to a LOB table space can change due to a recovery operation, depending on the type of recovery performed. If all of the following objects for all LOB columns are recovered in a single RECOVER utility statement to the present point-in-time, a QUIESCE point, or a COPY SHRLEVEL(REFERENCE) point, no pending status will exist:

- base table space
- index on the auxiliary table
- LOB table space

Refer to Table 54 on page 291 for information about the status of a base table space, index on the auxiliary table, or LOB table space that was recovered without its related objects.

Table 54. Determining object status after recovery.

Object	Recovery Type	Base table space Status	Index on the auxiliary table Status	LOB table space Status
Base table space	Current RBA or LRSN	None	None	None
Base table space	Point-in-time	CHECK-pending	None	None
Index on the auxiliary table	Current RBA or LRSN	None	None	None
Index on the auxiliary table	Point-in-time	None	CHECK-pending	None
LOB table space	Current RBA or LRSN, LOB table space defined with LOG(YES)	None	None	None
LOB table space	Current RBA or LRSN, LOB table space defined with LOG(NO)	None	None	auxiliary warning[1]
LOB table space	TOCOPY, COPY was SHRLEVEL REFERENCE	CHECK-pending	REBUILD-pending	None
LOB table space	TOCOPY, COPY was SHRLEVEL CHANGE	CHECK-pending	REBUILD-pending	CHECK-pending - auxiliary warning[1]
LOB table space	TOLOGPOINT or TORBA (not a quiesce point)	CHECK-pending	REBUILD-pending	CHECK-pending- auxiliary warning[1]
LOB table space	TOLOGPOINT or TORBA (at a quiesce point)	CHECK-pending	REBUILD-pending	None

Note: [1] If at any time a log record is applied to the LOB table space that results in a LOB being marked invalid, the LOB table space is set to auxiliary warning status.

For information about resetting any of these statuses, see “Appendix C. Resetting an advisory or restrictive status” on page 677.

Performing a point-in-time recovery

A recovery operation done with the options TOLOGPOINT, TORBA or TOCOPY is known as a point-in-time recovery. A consistent point-in-time is a quiesce point or an image copy set that was taken with SHRLEVEL REFERENCE. It is not necessary to take a full image copy after recovering to a point-in-time, except in the case of fallback recovery; see “Performing fallback recovery” on page 298. DB2 records the RBAs or LRSNs associated with the point-in-time recovery in the SYSIBM.SYSCOPY catalog table to allow future recover operations to skip the unwanted range of log records.

Because a point-in-time recovery of only the table space leaves data in a consistent state and indexes in an inconsistent state, all indexes must be rebuilt using REBUILD INDEX. For more information, see “Resetting the REBUILD-pending status” on page 269.

An index cannot be recovered to a prior point-in-time if the index has had its key length increased less than or equal to 16 distinct times since the specified point-in-time. An ALTER that increases the length of the index key column becomes distinct when the unit of work is committed— the ALTERs must take place in 16 different commit scopes. For example, if you ALTER, commit, and then ALTER again, this counts at two distinct ALTERs. Alternatively, if you ALTER, ALTER again and then commit, this counts as one distinct ALTER.

If you use TOLOGPOINT, TORBA or TOCOPY to recover a single data set of a nonpartitioned table space, DB2 issues message DSNU520I to warn that the table space can become inconsistent following the RECOVER job. This point-in-time recovery can cause compressed data to exist without a dictionary or can even overwrite the data set that contains the current dictionary.

The TORBA and TOLOGPOINT options reset the CHECK-pending status on table spaces when:

- All members of a table space set are recovered to the same quiesce point or SHRLEVEL REFERENCE point, and no referential constraints were defined on a dependent table after that point. The CHECK-pending status is reset for any table space in the table space set.

The TORBA and TOLOGPOINT options set the CHECK-pending status on table spaces when:

- One or more members of a table space set are recovered to a previous point in time that is not a common quiesce or SHRLEVEL(REFERENCE) point. Dependent table spaces are placed in CHECK-pending status.
- All members of a table space set are recovered to the same quiesce point, but referential constraints were defined on a dependent table after that quiesce point. Table spaces containing those dependent tables are placed in check pending status.
- Table spaces with LOB columns defined were recovered without recovering their LOB table spaces.

The TORBA and TOLOGPOINT options reset the CHECK-pending status on indexes when:

- The indexes have been recovered along with the related table space to the same quiesce point or SHRLEVEL REFERENCE point. RECOVER processing resets the CHECK-pending status for any indexes in the RECOVER statement.

The TORBA and TOLOGPOINT options set the CHECK-pending status on indexes when:

- One or more of the indexes have been recovered to a previous point in time, but the related table space was not recovered in the same RECOVER statement.
- One or more of the indexes have been recovered along with the table space to a previous point in time that is not a quiesce point or SHRLEVEL REFERENCE point.

The auxiliary check-pending status (ACHKP) is set on when the CHECK DATA utility detects an inconsistency between a base table space with LOB columns defined and a LOB table space. For information about how to reset the ACHKP status, see “Appendix C. Resetting an advisory or restrictive status” on page 677.

All user-defined table spaces and indexes with refresh-pending status (REFP) can be recovered with point-in-time recovery and either the TOCOPY, TORBA or TOLOGPOINT option.

For more information about recovering data to a prior point of consistency, see Part 4 (Volume 1) of *DB2 Administration Guide*.

Recovery considerations after rebalancing partitions with REORG: Image copies taken prior to resetting the REORG-pending status of any partition of a partitioned table space are not usable for recovering to a current RBA or LRSN.

Avoid performing a point-in-time recovery for a partitioned table space to a point-in-time that is after the REORG-pending status was set, but before a rebalancing REORG was performed. To determine an appropriate point-in-time:

1. Run REPORT RECOVERY.
2. Select an image copy where the recovery point is a point after the rebalancing REORG was performed.

If you run the REORG utility to turn off a REORG-pending status, and then recover to a point-in-time before that REORG, DB2 sets restrictive statuses on the all partitions that you specified in the REORG job as follows:

- Sets REORG-pending (and possibly CHECK-pending) on for the data partitions.
- Sets REBUILD-pending on for the associated index partitions.
- Sets REBUILD-pending on for the associated logical partitions of nonpartitioning indexes.

For information about resetting these restrictive statuses, see “REORG-pending status” on page 682 and “REBUILD-pending status” on page 680.

To create a new recovery point, it is strongly recommended that immediately following an ALTER INDEX operation, you either:

- Run REORG with COPYDDN and SHRLEVEL NONE specified, or
- Take a full image copy immediately after REORG completes.

Use RECOVER LOGONLY after data has been redistributed among partitions using REORG. If you perform a point-in-time recovery, you must keep the off-line copies synchronized with the SYSIBM.SYSCOPY records. Therefore, do *not* delete any SYSCOPY='A' records, as they might be needed during the recovery. Only delete these SYSCOPY records when you are sure you will no longer use the off-line copies taken before the rebalancing REORG.

Planning for point-in-time recovery: TOCOPY, TOLOGPOINT, and TORBA are viable alternatives in many situations in which recovery to the current point-in-time is not possible or desirable. To make these options work best for you, take periodic quiesce points at points of consistency that are appropriate to your applications.

When making copies of a single object, use SHRLEVEL REFERENCE to establish consistent points for TOCOPY recovery. Copies made with SHRLEVEL CHANGE do not copy data at a single instant, because changes can occur as the copy is made. A subsequent RECOVER TOCOPY operation can produce inconsistent data.

When copying a list of objects, use SHRLEVEL REFERENCE. If a subsequent recovery to a point-in-time is necessary, you can use a single RECOVER utility statement to list all of the objects, along with TOLOGPOINT to identify the common RBA or LRSN value. If you use SHRLEVEL CHANGE to copy a list of objects, you should follow it with a QUIESCE of the objects.

To improve the performance of the recovery, take a full image copy of the table space or table space set, and then quiesce them using the QUIESCE utility. This allows RECOVER TORBA to recover the table spaces to the quiesce point with minimal use of the log.

Authorization: Restrict use of TOCOPY, TOLOGPOINT, and TORBA to personnel with a thorough knowledge of the DB2 recovery environment.

Ensuring consistency: RECOVER TORBA, RECOVER TOLOGPOINT, and RECOVER TOCOPY can be used on a single:

- partition of a partitioned table space
- partition of a partitioning index space
- data set of a simple table space

For any of the previously-listed objects, all data sets must be restored to the same level or the data will be inconsistent.

If possible, a table space and all of its indexes (or a table space set and all related indexes) should be recovered in the same RECOVER utility statement, specifying TOLOGPOINT or TORBA to identify a QUIESCE point. This action avoids placing indexes in the CHECK-pending or REBUILD-pending status. If the TOLOGPOINT is not a common QUIESCE point for all objects, we recommend using the following procedure:

1. RECOVER table spaces to the log point.
2. Use concurrent REBUILD INDEX jobs to recover the indexes over each table space.

This procedure ensures that the table spaces and indexes are synchronized, and eliminates the need to run the CHECK INDEX utility.

Point-in-time recovery can cause table spaces to be placed in CHECK-pending status if they have table check constraints or referential constraints defined on them. When recovering tables involved in a referential constraint, you should recover all the table spaces involved in a constraint. This is the *table space set.* To avoid setting CHECK-pending, you must do both of the following:

- Recover the table space or the table space set to a quiesce point or to an image copy made with SHRLEVEL REFERENCE.

  If you do not recover each table space of the table space set to the same quiesce point, and if any of the table spaces are part of a referential integrity structure:

  – All dependent table spaces that are recovered are placed in CHECK-pending status with the scope of the whole table space.
  – All dependent table spaces of the above recovered table spaces are placed in CHECK-pending status with the scope of the specific dependent tables.
- Do not add table check constraints or referential constraints after the quiesce point or image copy.

  If you recover each table space of a table space set to the same quiesce point, but referential constraints were defined after the quiesce point, then the CHECK-pending status is set for the table space containing the table with the referential constraint.

For information about actions to take if CHECK INDEX identifies inconsistencies after you perform a RECOVER job, see “Reviewing CHECK INDEX output” on page 71.

For information about resetting the CHECK-pending status of table spaces, see “Chapter 6. CHECK DATA” on page 53. For information about resetting the CHECK-pending status for indexes, see “CHECK-pending status” on page 678.

Compressed data: Use caution when recovering a portion of a table space or partition, say one data set, to a prior point in time. If the data set being recovered

has been compressed with a different dictionary, then you can no longer read the data. The details of data compression are described in Part 5 (Volume 2) of *DB2 Administration Guide*.

Avoiding specific image copy data sets

You can accidentally lose an image copy, or you might want to avoid a specific image copy data set. Because the corresponding row is still present in SYSIBM.SYSCOPY, RECOVER will always attempt to allocate the data set. The following sections describe the options available if you want to skip a specific image copy data set.

Image copy on tape: If the image copy is on tape, message IEF233D and IEF455D will request the tape for RECOVER.

```
 IEF233D M BAB,COPY   ,,R92341QJ,DSNUPROC,
         OR RESPOND TO IEF455D MESSAGE
*42 IEF455D MOUNT COPY    ON BAB FOR R92341QJ,DSNUPROC OR REPLY 'NO'
 R 42,NO
 IEF234E K BAB,COPY    ,PVT,R92341QJ,DSNUPROC
```

By replying NO, you can initiate the fallback to the previous image copy. RECOVER will respond with message DSNU030I and DSNU508I.

```
DSNU030I    csect-name - UNABLE TO ALLOCATE R92341Q.UTQPS001.FCOPY010
                         RC=4, CODE=X'04840000'
DSNU508I    csect-name - IN FALLBACK PROCESSING TO PRIOR FULL IMAGE COPY
```

Reason code X'0484' means ″request denied by operator.″

Image copy on DASD: If the image copy is on DASD, you can delete or rename the image copy data set before RECOVER starts executing. RECOVER issues message DSNU030I and DSNU508I.

```
DSNU030I    csect-name - UNABLE TO ALLOCATE R92341Q.UTQPS001.FCOPY010,
                         RC=4, CODE=X'17080000'
DSNU508I    csect-name - IN FALLBACK PROCESSING TO PRIOR FULL IMAGE COPY
```

Reason Code X'1708' means ″ICF catalog entry not found″.

Improving performance

Use MERGECOPY to merge your table space image copies before recovering the table space. If you do not merge your image copies, RECOVER automatically merges them. If RECOVER cannot allocate all the incremental image copy data sets when it merges the image copies, then RECOVER uses the log instead.

Include a table space list in your RECOVER utility statement to avoid scanning the log more than once.

If you use RECOVER TOCOPY for full image copies, you can improve performance by using data compression. The improvement is proportional to the degree of compression.

Consider specifying the PARALLEL keyword to restore image copies from DASD to a list of objects in parallel.

Optimizing the LOGAPPLY phase

The time required to recover a table space depends also on the time required to read and apply log data. There are several things you can do to optimize the process.

If possible, the required log records are read from the *active* log. That provides the best performance.

Any log records not found in the active logs are read from the archive log data sets, which are dynamically allocated to satisfy the requests. The type of storage used for archive log data sets is a significant factor in the performance.

- RECOVER a list of objects in one utility statement to take only a single pass of the log.
- Keeping archive logs on DASD provides the best possible performance.
- Controlling archive log data sets by DFSMShsm is next best. DB2 optimizes recall of the data sets. After being recalled, the data set is read from DASD.
- If the archive log must be read from tape, DB2 optimizes access by means of *ready-to-process* and *look-ahead mount* requests. DB2 also permits delaying the deallocation of a tape drive if subsequent RECOVER jobs require the same archive log tape. Those methods are described in more detail below.

Which log data sets to use and where they reside is recorded in the BSDS, which must be kept current. If the archive log data sets are cataloged, the integrated catalog facility catalog tells where to allocate the required data set.

To improve recovery time, consider enabling the Fast Log Apply function on the DB2 subsystem. For more information about enabling this function, see the LOG APPLY STORAGE field on panel DSNTIPL, in Part 2 of *DB2 Installation Guide*.

DFSMShsm data sets: Recall for the first DFSMShsm archive log data set starts automatically when the LOGAPPLY phase starts. When recall is complete and the first log record is read, recall for the next archive log data set starts. This process is known as *look-ahead* recalling. Its object is to recall the next data set in parallel with reading the preceding one.

When a recall is complete, the data set is available to all RECOVER jobs that require it. Reading proceeds in parallel.

Non-DFSMShsm tape data sets: DB2 reports on the console all tape volumes that are required for the entire job. The report distinguishes two types of volumes:

- Any volume *not* marked with an asterisk (*) is *required* for the job to complete. These volumes should be obtained from the tape library as soon as possible.
- Any volume that *is* marked with an asterisk (*) contains data that is also contained in one of the active log data sets. The volume might or might not be required.

As tapes are mounted and read, DB2 makes two types of mount requests:

- *Ready-to-process*: The current job needs this tape immediately. As soon as the tape is loaded, DB2 allocates and opens it.
- *Look-ahead*: This is the next tape volume required by the current job. Responding to this request enables DB2 to allocate and open the data set before it is needed, thus reducing overall elapsed time for the job.

You can dynamically change the maximum number of input tape units that are used to read the archive log by specifying the COUNT option of the SET ARCHIVE command. For example, use

```
-SET ARCHIVE COUNT (10)
```

to assign 10 tape units to your DB2 subsystem.

The DISPLAY ARCHIVE READ command shows the currently mounted tape volumes and their statuses.

Delayed deallocation: DB2 can delay deallocating the tape units used to read the archive logs. This is useful when several RECOVER utility statements run in parallel. By delaying deallocation, DB2 can re-read the same volume on the same tape unit for different RECOVER jobs, without taking time to allocate it again.

The amount of time DB2 delays deallocation can be dynamically changed with the TIME option of the SET ARCHIVE command. For example, use:

```
-SET ARCHIVE TIME(60)
```

to specify a 60 minute delay. In a data sharing environment, you might want to specify (0) to avoid having one member hold onto a data set that another member needs for recovery.

Performance summary:

1. Achieve the best performance by allocating archive logs on DASD.
2. Consider staging cataloged tape data sets to DASD before allocation by the log read process.
3. If the data sets are read from tape, set both the COUNT and the TIME value to the maximum allowable within the system constraints.

Recovering image copies in a JES3 environment

Ensure that there are sufficient units available to mount the required image copies. In a JES3 environment, if the number of image copies to be restored exceeds the number of available online and offline units, and the RECOVER job successfully allocates all available units, the job waits for more units to become available.

Resetting RECOVER-pending or REBUILD-pending status

Several possible operations on a table space can place the table space in the RECOVER-pending status and the index space in REBUILD-pending status. The status can be turned off in several ways, listed below:

- Recover the table space, index space, or partition.
- Use REBUILD INDEX to rebuild the index space from existing data.
- Use the LOAD utility, with the REPLACE option, on the table space or partition.
- Use the REPAIR utility, with the NORCVRPEND option, on the table space, index space, or partition. Be aware that the REPAIR utility does not fix the data inconsistency in the table space or index.
- To rebuild indexes, run REORG TABLESPACE SORTDATA. Identify for table spaces and for indexes.

Considerations for running RECOVER

This section includes additional information to keep in mind when running RECOVER.

Allocating incremental image copies

RECOVER will attempt to dynamically allocate ALL required incremental image copy data sets. If any of the incremental image copies are missing, RECOVER will:

- Identify the first incremental image copy that is missing
- Use the incremental image copies up to the missing incremental image copy
- Not use the remaining incremental image copy data sets

- Apply additional log records to compensate for any incremental image copies that were not used

For example, if the incremental image copies are on tape and not enough tape drives are available, RECOVER will NOT use the remaining incremental image copy data sets.

Performing fallback recovery

If the RECOVER utility cannot use the latest primary copied data set as a starting point for recovery, it attempts to use the backup copied data set, if one is available. If neither image copy is usable, the RECOVER utility attempts to fall back to a previous recoverable point. If the previous recoverable point is a full image copy, the RECOVER utility uses the full image copy, any incremental image copies, and the log to recover. If a previous REORG LOG YES or LOAD REPLACE LOG YES was done, the RECOVER utility attempts to recover from the log and applies any changes that occurred between the two image copies. If there are no good full image copies, and no previous REORG LOG YES or LOAD REPLACE LOG YES, the RECOVER utility terminates.

If RECOVER processes an index for which no full copy exists, or if the copy cannot be used due to utility activity that occurred on the index or on its underlying table space, the index remains untouched and utility processing terminates with return code 8. For more information about this situation, see “Setting and clearing the informational COPY-pending status” on page 100.

If you always make multiple image copies, RECOVER should seldom fall back to an earlier point. Instead, RECOVER relies on the backup copied data set should the primary copied data set be unusable.

In a JES3 environment, fallback recovery can be accomplished by issuing a JES3 ″cancel,s″ command at the time the allocation mount message is issued. This might be necessary in the case where a volume is not available or the given volume is not desired.

Retaining tape mounts

If the image copy data sets from which you want to recover reside on the same tape, you do not need to remove the tape. For noncataloged image copies, specify the following parameters on the DD cards (in this example, the DD cards are *ddname1* and *ddname2*):

```
//ddname1   DD   UNIT=TAPE,DSN=data-set-name1,DISP=(OLD,PASS),LABEL=1,
//               VOL=(,RETAIN,SER=vol-ser)
//ddname2   DD   UNIT=TAPE,DSN=data-set-name2,DISP=(OLD,PASS),LABEL=2,
//               VOL=(,REF=*.ddname1)
```

This example only works for multiple image copies on a single volume. To use multiple image copies on multiple volumes, the image copy data sets must be cataloged. For cataloged image copies on one or more tape volumes, specify the following parameters on the DD cards (in this example, the DD cards are *ddname1*, *ddname2*, and *ddname3*):

```
//ddname1   DD   DSN=data-set-name1,UNIT=3480,DISP=(OLD,PASS),VOL=(,RETAIN),
//               LABEL=(1,SL)
//ddname2   DD   DSN=data-set-name2,UNIT=3480,DISP=(OLD,PASS),VOL=(,RETAIN),
//               LABEL=(2,SL)
//ddname3   DD   DSN=data-set-name3,UNIT=3480,DISP=(OLD,PASS),VOL=(,RETAIN),
//               LABEL=(3,SL)
```

Avoiding damaged media

When a media error is detected, DB2 prints a message giving the extent of the damage. If an entire volume is bad, and storage groups are being used, you must remove the bad volume first; otherwise, the RECOVER utility can re-access the damaged media. You must:

1. Use -ALTER STOGROUP to remove the bad volume and add another
2. Execute the RECOVER utility for all objects on that volume.

If the RECOVER utility cannot complete because of severe errors caused by the damaged media, it can be necessary to use access method services (IDCAMS) to delete the cluster for the table space or index with the NOSCRATCH option. Refer to the access method services reference manual for details. If the table space or index is defined using STOGROUP, the RECOVER utility automatically redefines the cluster. For user-defined table spaces or indexes, you must redefine the cluster before invoking the RECOVER utility.

Recovering table spaces and index spaces with mixed volume IDs

You cannot run RECOVER on a table space or index space on which mixed specific and non-specific volume IDs were defined with CREATE STOGROUP or ALTER STOGROUP.

Terminating or restarting RECOVER

For instructions on restarting a utility job, see “Restarting an online utility” on page 46 .

Terminating RECOVER: Terminating a RECOVER job with the -TERM UTILITY command leaves the table space being recovered in RECOVER-pending status, and the index space being recovered in the REBUILD pending status. If you are recovering a table space to a previous point-in-time, then its indexes are left in the REBUILD-pending status. The data is unavailable until the object has been successfully recovered or rebuilt.

Restarting RECOVER: RECOVER can be restarted at the beginning of the phase or at the current checkpoint.

If you attempt to recover multiple objects using a single RECOVER statement and the utility fails in:

- The RESTORE phase: All objects in the process of being restored are placed in the RECOVER or REBUILD-pending status. The status of the remaining objects is unchanged.
- The LOGAPPLY phase: All objects specified in the RECOVER statement are placed in the RECOVER or REBUILD-pending status.

In both cases, you must identify and fix the causes of the failure before a current restart is performed.

If you specified the PARALLEL keyword in your RECOVER utility statement, use RESTART(PHASE) to restart at the beginning of the current phase, or restart from the last commit point of each object processed in parallel.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible. However, if a nonpartitioning index exists on a partitioned table space, utilities operating on different partitions of a table space can be incompatible because of contention on the nonpartitioning index.

Table 55 shows which claim classes RECOVER claims and drains and any restrictive state the utility sets on the target object.

Table 55. Claim classes of RECOVER operations. Use of claims and drains; restrictive states set on the target object.

Target	RECOVER (no option)	RECOVER TORBA or TOCOPY	RECOVER PART TORBA or TOCOPY	RECOVER ERROR-RANGE
Table space or partition	DA/UTUT	DA/UTUT	DA/UTUT	DA/UTUT CW/UTRW[1]
Index or physical partition	DA/UTUT	DA/UTUT	DA/UTUT	DA/UTUT CW/UTRW[1]
Nonpartitioning index	DA/UTUT	DA/UTUT		DA/UTUT CW/UTRW[1]
RI dependents		CHKP (YES)	CHKP (YES)	

Legend:
- CHKP (YES): Concurrently running applications see CHECK-pending after commit
- CW: Claim the write claim class
- DA: Drain all claim classes, no concurrent SQL access
- DR: Drain the repeatable read class, no concurrent access for SQL repeatable readers
- UTRW: Utility restrictive state, read/write access allowed
- UTUT: Utility restrictive state, exclusive control
- Blank: Object is not affected by this utility

Notes:

1. During the UTILINIT phase, the claim and restrictive state change from DA/UTUT to CW/UTRW.

RECOVER does not set a utility restrictive state if the target object is DSNDB01.SYSUTILX.

Table 56 shows which utilities can run concurrently with RECOVER on the same target object. The target object can be a table space, an index space, or a partition of a table space or index space. If compatibility depends on particular options of a utility, that is also shown.

Table 56. RECOVER compatibility

Action	RECOVER (no option)	RECOVER TOCOPY or TORBA	RECOVER ERROR-RANGE
CHECK DATA	No	No	No
CHECK INDEX	No	No	No
CHECK LOB	No	No	No
COPY INDEXSPACE	No	No	No

Table 56. RECOVER compatibility (continued)

Action	RECOVER (no option)	RECOVER TOCOPY or TORBA	RECOVER ERROR-RANGE
COPY TABLESPACE	No	No	No
DIAGNOSE	Yes	Yes	Yes
LOAD	No	No	No
MERGECOPY	No	No	No
MODIFY	No	No	No
QUIESCE	No	No	No
REBUILD INDEX	No	No	No
REORG INDEX	Yes	No	Yes
REORG TABLESPACE	No	No	No
REPAIR LOCATE TABLESPACE	No	No	No
REPAIR LOCATE INDEX	Yes	No	Yes
REPORT	Yes	Yes	Yes
RUNSTATS INDEX	No	No	No
RUNSTATS TABLESPACE	No	No	No
STOSPACE	Yes	Yes	Yes
UNLOAD	No	No	No

To run on DSNDB01.SYSUTILX, RECOVER must be the only utility in the job step and the only utility running in the DB2 subsystem.

RECOVER on any catalog or directory table space is an “exclusive” job; such a job can interrupt another job between job steps, possibly causing the interrupted job to time out.

Sample control statements

Example 1: Recover an error range. Recover from reported media failure in partition 2 of table space DSN8D71A.DSN8S71D.

```
//STEP5    EXEC DSNUPROC,UID='HUIAU326.RESTORE',TIME=1440,
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//SYSIN    DD *
RECOVER TABLESPACE DSN8D71A.DSN8S71D DSNUM 2 ERROR RANGE
/*
```

Example 2: Recover a table space. Recover table space DSN8S71D, in database DSN8D71A.

```
RECOVER TABLESPACE DSN8D71A.DSN8S71D
```

Example 3: Recover a table space partition. Recover the second partition of table space DSN8S71D.

```
RECOVER TABLESPACE DSN8D71A.DSN8S71D DSNUM 2
```

Example 4: Recover a table space to a specific RBA. Recover table spaces DSN8D71A.DSN8S71E and DSN8D71A.DSN8S71D to their quiesce point (RBA `X'000007425468'`).

```
RECOVER TABLESPACE DSN8D71A.DSN8S71E DSNUM 2
        TABLESPACE DSN8D71A.DSN8S71D
        TORBA X'000007425468'
```

Example 5: Recover a list of objects to a point in time. The point in time is the common LRSN value from the SYSIBM.SYSCOPY records for the list of objects in the COPY SHRLEVEL REFERENCE job on page 105. The objects in the list are synchronized after successful completion of this RECOVER utility statement. This example restores four objects in parallel.

```
RECOVER TOLOGPOINT X'1234567890AB' PARALLEL(4)
      TABLESPACE DSN8D71A.DSN8S71D
      INDEX DSN8710.XDEPT1
      INDEX DSN8710.XDEPT2
      INDEX DSN8710.XDEPT3
      TABLESPACE DSN8D71A.DSN8S71E
      INDEX DSN8710.XEMP1
      INDEX DSN8710.XEMP2
```

Example 6: Recover an index space partition to the last image copy taken. Recover the first partition of index space IPDH082P in database DBDH082P to the last image copy taken. If the last image copy taken is a full image copy, then this full image copy is restored. If it is an incremental image copy, then the most recent full image copy along with any incremental image copies are restored.

```
RECOVER INDEXSPACE DBDH082P.IPDH082P DSNUM 1 TOLASTCOPY
```

Example 7: Recover an index to the last full image copy taken without deleting and refining the data sets. Any incremental image copies that were taken after the full image copy are not used to restore the object. REUSE specifies that DB2 logically resets and reuses DB2-managed data sets without deleting and redefining them.

```
RECOVER INDEX       ADMF001.IADH082P REUSE TOLASTFULLCOPY
```

Chapter 22. REORG INDEX

The REORG INDEX utility reorganizes an index space to improve access performance and reclaim fragmented space. You can specify the degree of access to your data during reorganization, and collect inline statistics using the STATISTICS keyword.

You can determine when to run REORG INDEX by using the LEAFDISTLIMIT catalog query option. If you specify the REPORTONLY option, REORG INDEX will produce a report detailing if a REORG is recommended; a REORG is not performed. These options are not available for indexes on the directory.

For a diagram of REORG INDEX syntax and a description of available options, see "Syntax and options of the control statement" on page 304. For detailed guidance on running this utility, see "Instructions for running REORG INDEX" on page 314.

Output: The following list summarizes REORG INDEX output:

REORG specified	Results
REORG INDEX	Reorganizes the entire index (all parts if partitioning).
REORG INDEX PART *n*	Reorganizes PART *n* of the partitioning index.

Authorization required: To execute this utility on a user index, the privilege set of the process must include one of the following:
- REORG privilege for the database
- DBADM or DBCTRL authority for the database
- SYSCTRL authority
- SYSADM authority.

To execute this utility on an index space in the catalog or directory, the privilege set of the process must include one of the following:
- REORG privilege for the DSNDB06 (catalog) database
- DBADM or DBCTRL authority for the DSNDB06 (catalog) database
- Installation SYSOPR authority
- SYSCTRL authority
- SYSADM or Installation SYSADM authority

An authority other than installation SYSADM or installation SYSOPR can receive message DSNT500I, "resource unavailable", while trying to reorganize an index space in the catalog or directory. This can happen when the DSNDB06.SYSDBAUT or DSNDB06.SYSUSER catalog table space or one of the indexes is unavailable. If this problem occurs, run the REORG INDEX utility again using an authorization ID with the installation SYSADM or installation SYSOPR authority.

An ID with installation SYSOPR authority can also execute REORG INDEX, but only on an index in the DSNDB06 database.

To run REORG INDEX STATISTICS REPORT YES, the privilege set must include the SELECT privilege on the catalog tables.

Execution phases of REORG INDEX: The REORG utility operates in these phases:

Phase	Description

© Copyright IBM Corp. 1983, 2001

UTILINIT	Initialization and setup
UNLOAD	Unloads index space and writes keys to a sequential data set.
BUILD	Builds indexes. Updates index statistics.
LOG	Processes log iteratively. Used only if you specify SHRLEVEL CHANGE.
SWITCH	Switches access between original and new copy of index space or partition. Used only if you specify SHRLEVEL REFERENCE or CHANGE.
UTILTERM	Cleanup. For DB2-managed data sets and either SHRLEVEL CHANGE or SHRLEVEL REFERENCE, the utility deletes the original copy of the table space or index space.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

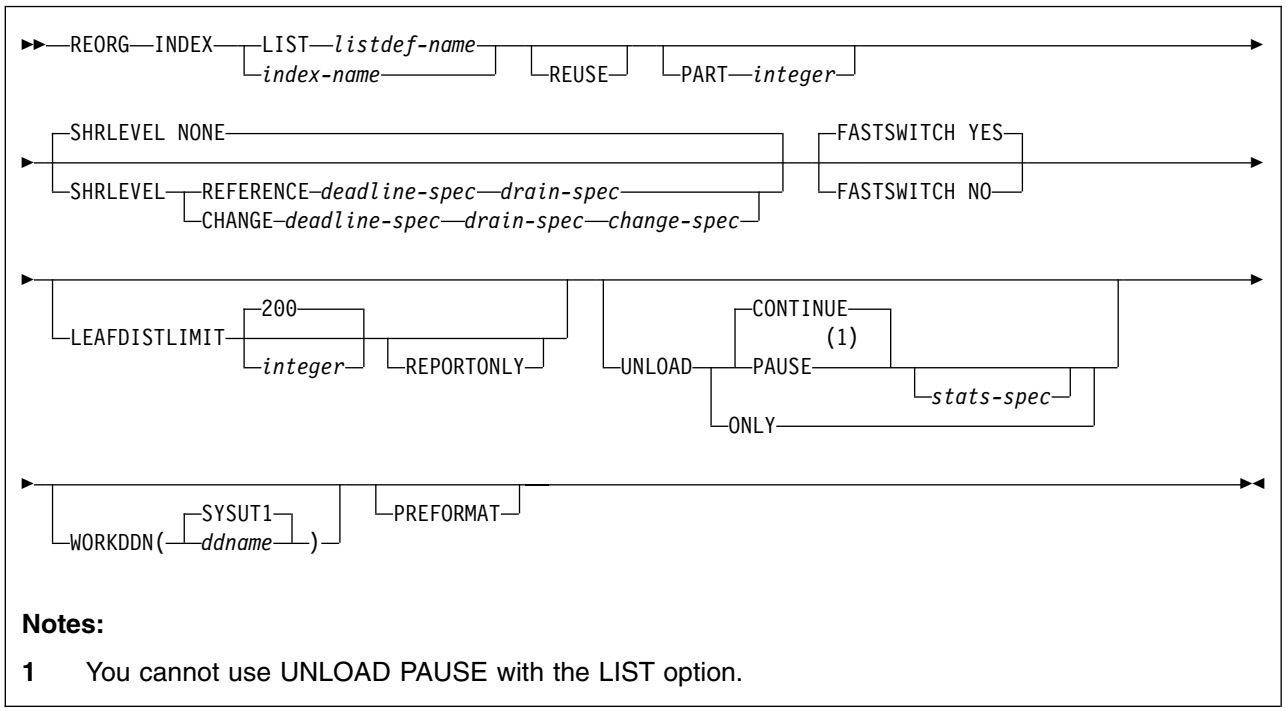

Notes:

1 You cannot use UNLOAD PAUSE with the LIST option.

deadline-spec:

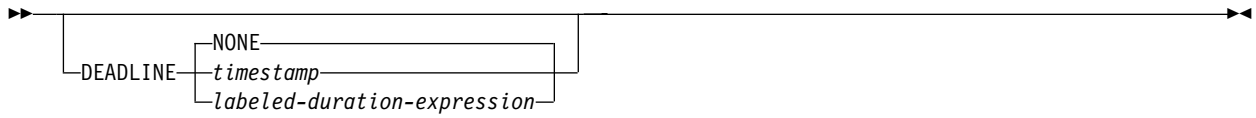

drain-spec:

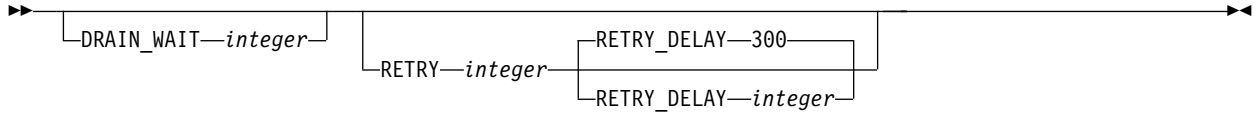

change-spec:

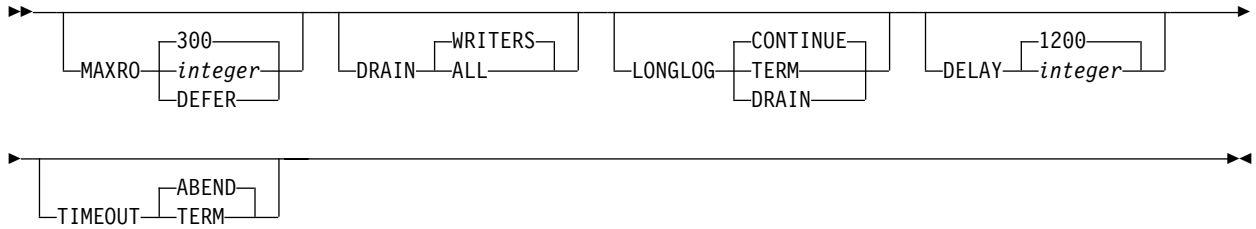

labeled-duration-expression:

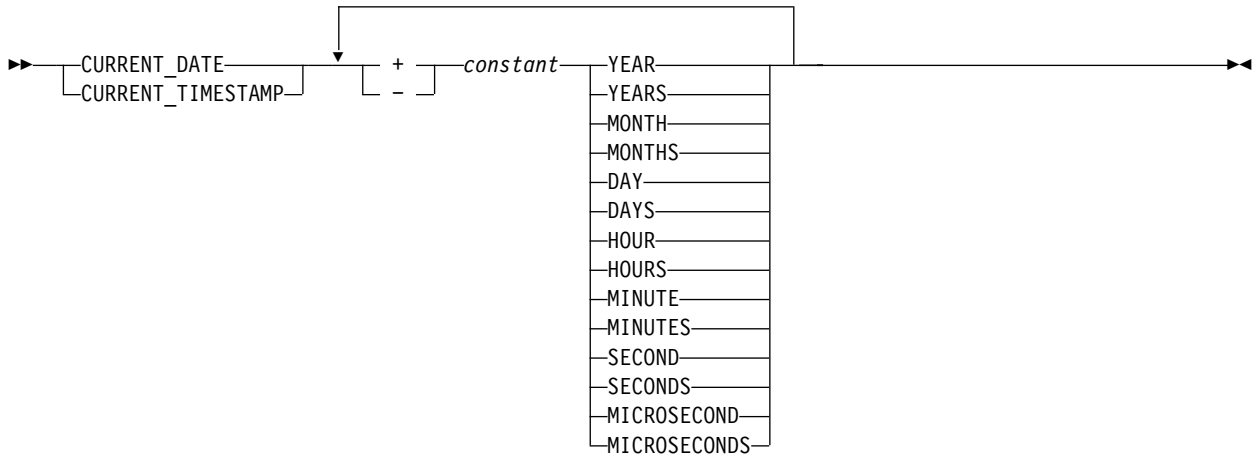

stats-spec:

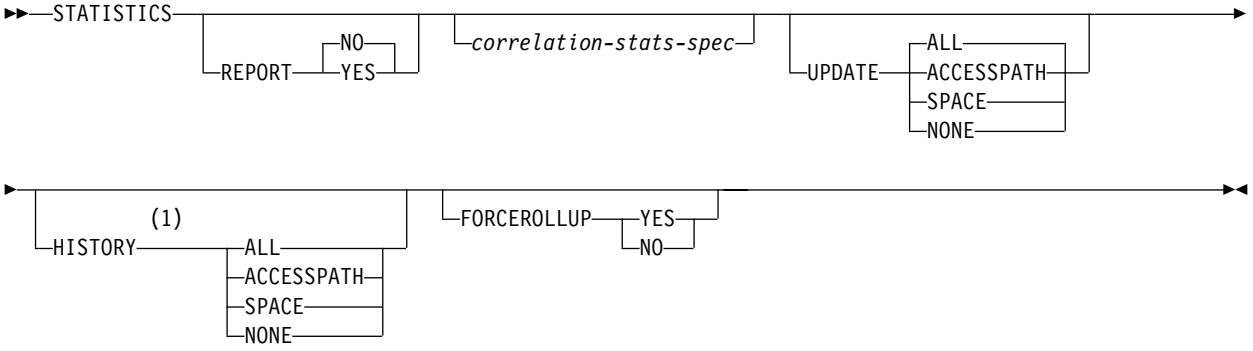

Notes:

1 What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

correlation-stats-spec:

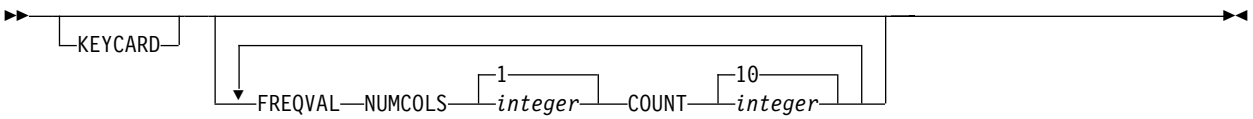

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

INDEX *index-name*
Specifies an index to be reorganized.

index-name is the qualified name of the index, in the form *creator-id.index-name*. If you omit the qualifier creator ID, the user identifier for the utility job is used.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name. The INDEX keyword is required to differentiate this REORG INDEX LIST from REORG TABLESPACE LIST. The utility allows one LIST keyword for each control statement of REORG INDEX. The list must contain only index spaces. Utility list manager will invoke REORG INDEX once for each item in the list. For more information about LISTDEF specifications, see “Generating object lists using LISTDEF” on page 144.

Do not specify STATISTICS INDEX *index-name* with REORG INDEX LIST. If you want to collect inline statistics for a LIST of indexes, specify STATISTICS INDEX (ALL).

As a rule, DSNUM and PART can never be specified with LIST on any utility.

REUSE
When used with SHRLEVEL NONE, specifies that REORG logically resets and reuses DB2-managed data sets without deleting and redefining them. If you do not specify REUSE and SHRLEVEL NONE, DB2 deletes and redefines DB2-managed data sets to reset them.

If a data set has multiple extents, the extents will not be released if you use the REUSE parameter.

REUSE does not apply if you also specified SHRLEVEL REFERENCE or CHANGE.

PART *integer*
Identifies a partition to be reorganized. You can reorganize a single partition of a partitioning index. PART can not be specified with LIST. *integer* must be in the range from 1 to the number of partitions that are defined for the partitioning index. The maximum is 254.

integer Designates a single partition.

If you omit the PART keyword, the entire index is reorganized.

SHRLEVEL
Specifies the method for performing the reorganization. The parameter following SHRLEVEL indicates the type of access allowed during the RELOAD phase of REORG.

NONE Specifies that reorganization operates by unloading from the area being reorganized (while applications can read but cannot write to the area), building into that area (while applications have no access), and then allowing read/write access again.

The **default** is **NONE**.

If you specify NONE (explicitly or by default), the following parameters cannot be specified: MAXRO, LONGLOG, DELAY, and DEADLINE.

REFERENCE
Specifies that reorganization operates by unloading from the area being reorganized (while applications can read but cannot write to the area), building into a shadow copy of that area (while applications can read but cannot write to the original copy), switching applications' future access from the original copy to the shadow copy by exchanging the names of the data sets, and then allowing read/write access again.

To determine which data sets are required when you execute REORG SHRLEVEL REFERENCE, see "Data sets used by REORG INDEX" on page 316.

If you specify REFERENCE, you cannot specify the following parameters:

- UNLOAD. Reorganization with REFERENCE always performs UNLOAD CONTINUE.
- MAXRO, LONGLOG, and DELAY.

CHANGE
Specifies that reorganization operates by unloading from the area being reorganized (while applications can read and write to the area), building into a shadow copy of that area (while applications have read/write access to the original copy of the area), applying the log of the original copy to the shadow copy (while applications can read and usually write

to the original copy), switching applications' future access from the original copy to the shadow copy by exchanging the names of the data sets, and then allowing read/write access again.

To determine which data sets are required when you execute REORG SHRLEVEL CHANGE, see "Data sets used by REORG INDEX" on page 316.

If you specify CHANGE, you cannot specify the following parameters:

- UNLOAD. Reorganization with CHANGE always performs UNLOAD CONTINUE.

FASTSWITCH *integer*

The **default** FASTSWITCH **YES** enables the SWITCH phase methodology. You have a choice between which methodology to use for a given reorganization. FASTSWITCH NO will cause the SWITCH phase to use IDCAMS RENAME.

When using the inline COPY function, the FASTSWITCH methodology is used in the SWITCH phase. The default for the FASTSWITCH parameter is YES. If you do not specify FASTSWITCH NO, the fifth level qualifier in the data set name alternates between I0001 and J0001.

DEADLINE

Specifies the deadline for the switch phase to finish. If DB2 estimates that the SWITCH phase will not finish by the deadline, DB2 issues the messages that the -DISPLAY UTILITY command would issue and then terminates reorganization.

NONE

Specifies that there is no deadline by which the SWITCH phase of log processing must finish.

The **default** is **NONE.**

timestamp

timestamp specifies the deadline for the SWITCH phase of log processing to finish. This deadline must not have already occurred when REORG is executed.

labeled-duration-expression

Calculates the deadline for the SWITCH phase of log processing to finish. The calculation is either based on CURRENT TIMESTAMP or CURRENT DATE. This deadline must not have already occurred when REORG is executed.

For example, to ensure that the SWITCH phase is complete by 6:30 AM two days from now, use the following expression:

```
DEADLINE CURRENT DATE + 2 DAYS + 6 HOURS + 30 MINUTES
```

If REORG SHRLEVEL REFERENCE or SHRLEVEL CHANGE terminates due to a DEADLINE specification, DB2 issues message DSNU374I with reason code 2, but does not set a restrictive status.

DRAIN_WAIT *integer*

Specifies the number of seconds that the utility will wait for DML (inserts, updates, deletes and select / fetches) when draining. The time specified is the aggregate time for all partitions of the index to be reorganised. This overrides the values specified by IRLMRWT and UTIMOUT, for DML only. For other tasks such as COMMANDS, the IRLMRWT and UTIMOUT values are used. If the

keyword is omitted or a value of 0 is specified, then regular draining using the IRLMRWT and UTIMOUT values will be used. Acceptable values can be from 0 to 1800.

RETRY *integer*
Specifies the maximum number of retries that can be attempted. Values can be from 0 to 255. If the keyword is omitted, then no retries will be attempted.

Note: Specifying RETRY can lead to increased processing costs and can result in multiple or extended periods of read-only access.

RETRY_DELAY
Specifies the minimum duration in seconds between retries. Values can be from 1 to 1800.

The **default** is **300** seconds. The value must be an integer.

MAXRO *integer*
Specifies the maximum amount of time for the last iteration of log processing. During that iteration, applications have read-only access.

The actual execution time of the last iteration might exceed the value specified with MAXRO.

The ALTER UTILITY command can change the value of MAXRO.

The **default** is **300** seconds. The value must be an integer.

integer
integer is the number of seconds. Specifying a small positive value reduces the length of the period of read-only access, but it might increase the elapsed time for REORG to complete. If you specify a huge positive value, the second iteration of log processing is probably the last iteration.

DEFER
Specifies that the iterations of log processing with read/write access can continue indefinitely. REORG never begins the final iteration with read-only access, unless you change the MAXRO value with ALTER UTILITY.

If you specify DEFER, you should also specify LONGLOG CONTINUE.

If you specify DEFER, and DB2 determines that the actual time for an iteration and the estimated time for the next iteration are both less than 5 seconds, DB2 adds a 5 second pause to the next iteration. This pause reduces consumption of processor time. The first time this situation occurs for a given execution of REORG, DB2 sends message DSNU362I to the console. The message states that the number of log records that must be processed is small and that the pause will occur. The message also suggests that this would be an appropriate time to execute ALTER UTILITY to change the MAXRO value and thus cause REORG to finish. DB2 adds the pause whenever the situation occurs; however, DB2 sends the message only if 30 minutes have elapsed since the last message was sent for a given execution of REORG.

DRAIN
Specifies drain behavior at the end of the log phase after the MAXRO threshold is reached and when the last iteration of the log is to be applied.

WRITERS
Specifies the current default action, in which DB2 drains just writers

during the log phase after the MAXRO threshold is reached and subsequently issues DRAIN ALL on entering the switch phase.

ALL Specifies that DB2 drain all readers and writers during the log phase, after the MAXRO threshold is reached.

Consider specifying DRAIN ALL if the following conditions are both true:

- There is a lot of SQL update activity during the log phase.
- The default behavior results in a large number of -911 SQL error messages.

LONGLOG

Specifies the action that DB2 performs, after sending a message to the console, if the number of records that the next iteration of log process will process is not sufficiently lower than the number that the previous iterations processed. This situation means that reorganization's reading of the log is not catching up to applications' writing of the log quickly enough.

CONTINUE

Specifies that until the time on the JOB statement expires, DB2 continues performing reorganization, including iterations of log processing, if the estimated time to perform an iteration exceeds the time specified with MAXRO.

A value of DEFER for MAXRO and a value of CONTINUE for LONGLOG together mean that REORG continues allowing access to the original copy of the area being reorganized and does not switch to the shadow copy. The user can execute the -ALTER UTILITY command with a large value for MAXRO when the switching is desired.

The **default** is **CONTINUE.**

TERM Specifies that DB2 will terminate reorganization after the delay specified by the DELAY parameter.

DRAIN

Specifies that DB2 drains the write claim class after the delay specified by the DELAY parameter. This action forces the final iteration of log processing to occur.

DELAY *integer*

Specifies the minimum interval between the time that REORG sends the LONGLOG message to the console and the time REORG that performs the action specified by the LONGLOG parameter.

integer is the number of seconds. The value must be an integer. The **default** is **1200**.

TIMEOUT

Specifies the action to be taken if the REORG utility gets a time out condition while trying to drain objects in either the LOG or SWITCH phases.

ABEND

If a time out condition occurs, DB2 leaves the objects in a UTRO or UTUT state.

TERM

If you specify the TERM option and a time out condition occurs, then DB2:

1. Issues an implicit TERM UTILITY command, causing the utility to end with a return code 8
2. Issues the DSNU590I and DSNU170I messages

3. Leaves the objects in a RW state.

LEAFDISTLIMIT
The specified value is compared to LEAFDIST for the specified partitions in SYSIBM.SYSINDEXPART for the specified index. If any LEAFDIST exceeds the value specified for LEAFDISTLIMIT, REORG is performed or recommended.

The **default** value is **200**.

REPORTONLY
If you specify this option, the REORG is only recommended, not performed. REORG produces a report with one of the following return codes:

1 No limit met; no REORG performed or recommended.
2 REORG performed or recommended.

UNLOAD
Specifies whether the utility job is to continue processing or end after the data is unloaded.

CONTINUE
Specifies that, after the data has been unloaded, the utility continues processing.

The **default** is **CONTINUE**.

PAUSE
Specifies that after the data has been unloaded, processing ends. The utility stops and the RELOAD status is stored in SYSIBM.SYSUTIL so that processing can be restarted with RELOAD RESTART(PHASE).

This option is useful if you want to redefine data sets during reorganization. For example, with a user defined data set, you can:
- Run REORG with the UNLOAD PAUSE option
- Redefine the data set using access method services
- Restart REORG by resubmitting the previous job and specifying RESTART(PHASE).

| If zero records are unloaded during an UNLOAD PAUSE, when
| REORG is restarted, the RELOAD and BUILD phases are bypassed.

| However, you cannot use UNLOAD PAUSE if you specify the LIST
| option.

ONLY Specifies that after the data has been unloaded, the utility job ends and the status in SYSIBM.SYSUTIL corresponding to this utility ID is removed.

STATISTICS
Specifies the gathering of statistics for the index; the statistics are either
reported or stored in the DB2 catalog. You cannot collect inline statistics for
indexes on the catalog and directory tables.

REPORT
Determines if a set of messages is generated to report the collected statistics.

NO
Indicates that the set of messages is not output to SYSPRINT.

The **default** is **NO**.

YES
Indicates that the set of messages is output to SYSPRINT. The messages generated are dependent on the combination of keywords (such as

TABLESPACE, INDEX, TABLE, and COLUMN) specified with the RUNSTATS utility. However, these messages are *not* dependent on the specification of the UPDATE option. REPORT YES always generates a report of SPACE and ACCESSPATH statistics.

UPDATE
Tells whether the collected statistics are inserted into the catalog tables. UPDATE also allows you to select statistics used for access path selection or statistics used by database administrators. What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

ALL Indicates that all collected statistics will be updated in the catalog.

The **default** is **ALL**.

ACCESSPATH
Indicates that only the catalog table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only the catalog table columns that provide statistics to help the database administrator assess the status of a particular table space or index are updated.

NONE Indicates that no catalog tables are updated with the collected statistics. This option is only valid when REPORT YES is specified.

HISTORY
Records all catalog table inserts or updates to the catalog history tables. What you specify under the UPDATE option affects what options you can specify under the HISTORY option. The table below lists the HISTORY options allowed for each UPDATE option specified.

The default will be supplied by the value specified in STATISTICS HISTORY on panel DSNTIPO.

ALL Indicates that all collected statistics will be updated in the catalog history tables.

ACCESSPATH
Indicates that only the catalog history table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only SPACE-related catalog statistics are updated in catalog history tables.

NONE Indicates that no catalog history tables are updated with the collected statistics.

Table 57. Allowable HISTORY options based on selected UPDATE option

UPDATE option	HISTORY options
ALL	ALL, ACCESSPATH, SPACE, NONE
ACCESSPATH	ACCESSPATH, NONE
SPACE	SPACE, NONE
NONE	NONE

FORCEROLLUP
Determines if aggregation or rollup of statistics will take place when RUNSTATS is executed in spite of some parts being empty. This will enable the optimizer to select the best access path.

The following are the available options for the FORCEROLLUP Keyword.

YES Indicates that forced aggregation or rollup processing is to be done, even though some parts may not contain data.

NO Indicates that aggregation or rollup will be done only if data is available for all parts.

If data is not available for all parts, DSNU623I message will be issued, provided the installation value for STATISTICS ROLLUP on panel DSNTIPO is set to NO.

KEYCARD
Collects all of the distinct values in all of the 1 to *n* key column combinations for the specified indexes. *n* is the number of columns in the index.

FREQVAL
Controls the collection of frequent value statistics. If you specify FREQVAL, it must be followed by two additional keywords:

NUMCOLS
Indicates the number of key columns to concatenate together when collecting frequent values from the specified index. Specifying '3' means to collect frequent values on the concatenation of the first three key columns. The **default** is **1**, which means collect frequent values on the first key column of the index.

COUNT
Indicates the number of frequent values to be collected. Specifying '15' means collect 15 frequent values from the specified key columns. The **default** is **10**.

WORKDDN(*ddname***)**
ddname specifies the DD statement for the unload data set.

ddname
Is the DD name of the temporary work file for build input.

The **default** is **SYSUT1**.

The WORKDDN keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

Even though WORKDDN is an optional keyword, a DD card for the unload output data set is required in the JCL. If you do not specify WORKDDN, or if you specify it without a *ddname*, the JCL must have a DD card with the name SYSUT1. If *ddname* is given, then a DD card or TEMPLATE must be supplied that matches it.

PREFORMAT
Specifies that the remaining pages are preformatted up to the high allocated RBA in the index space. The preformatting occurs after the index is built.

PREFORMAT can operate on an entire index space, or on a partition of a partitioned index space.

PREFORMAT is ignored if you specify UNLOAD ONLY.

For more information on PREFORMAT, see “Improving performance with LOAD or REORG PREFORMAT” on page 197.

Instructions for running REORG INDEX

To run REORG INDEX, you must:

1. Read “Before running REORG INDEX” in this chapter.
2. Prepare the necessary data sets, as described in “Data sets used by REORG INDEX” on page 316.
3. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for REORG INDEX, see “Sample control statements” on page 325.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 316. (For a complete description of the syntax and options for REORG INDEX, see “Syntax and options of the control statement” on page 304.)
5. Check the compatibility table in “Concurrency and compatibility” on page 322 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the REORG job doesn’t complete, as described in “Terminating or restarting REORG INDEX” on page 321.
7. Run REORG.

See “Chapter 3. Invoking DB2 online utilities” on page 23 for an explanation of ways to execute DB2 utilities.

Before running REORG INDEX

Region size: The recommended minimum region size is 4096K.

User-managed data sets and SHRLEVEL REFERENCE and CHANGE: When you execute the REORG utility with SHRLEVEL REFERENCE or SHRLEVEL CHANGE, the shadow data sets already exist. The name of the shadow data set is catname.DSNDBx.dbname.psname.y0001.Annn, where y is I or J.

To determine what the existing value is, execute the appropriate query against SYSTABLEPART or SYSINDEXPART:

```
SELECT DBNAME, TSNAME, IPREFIX
  FROM SYSIBM.SYSTABLEPART
  WHERE DBNAME = 'dbname' AND TSNAME = 'psname';
SELECT DBNAME, IXNAME, IPREFIX
  FROM SYSIBM.SYSINDEXES X, SYSIBM.SYSINDEXPART Y
  WHERE X.NAME = Y.IXNAME AND X.CREATOR = Y.IXCREATOR
  AND X.DBNAME = 'dbname' AND X.INDEXSPACE = 'psname';
```

With a partitioned table space, you get several rows back, and you need to find the row for the partitions that you want to reorganize.

The names have the form `catname.DSNDBx.dbname.psname.x0001.Annn`. Define the data sets as LINEAR and use SHAREOPTIONS(3,3). An efficient method for defining shadow data sets specifies MODEL, so that the utility creates the shadow like the original. For example:

```
DEFINE CLUSTER +
       (NAME('catname.DSNDBC.dbname.psname.x0001.A001') +
       MODEL('catname.DSNDBC.dbname.psname.y0001.A001')) +
```

```
DATA                                                +
 (NAME('catname.DSNDBD.dbname.psname.x0001.A001') +
 MODEL('catname.DSNDBD.dbname.psname.y0001.A001') )
```

If an index or partition index resides in DB2-managed data sets and shadow data sets do not already exist when you execute REORG, DB2 creates the shadow data sets. At the end of REORG processing, the DB2-managed shadow data sets are deleted. You can create the shadows ahead of time for DB2-managed data sets, and it is strongly recommended that you do so for the shadow data set of the logical partition of nonpartitioning indexes.

Estimating the size of pre-allocated data sets: If you have not changed the value of FREEPAGE or PCTFREE, the amount of space required for a shadow data set should be comparable to the amount of space required for the original data set.

Restart-pending status and SHRLEVEL CHANGE: If you specify SHRLEVEL CHANGE, REORG drains the write claim class near the end of REORG processing. In a data sharing environment, if a data sharing member fails and that member has restart pending status for a target page set, the drain can fail. You must postpone running REORG with SHRLEVEL CHANGE until all restart pending statuses have been removed. You can use the DISPLAY GROUP command to determine whether a member's status is FAILED. You can use the DISPLAY DATABASE command with the LOCKS option to determine if locks held.

Data sharing considerations for REORG: You must not execute REORG on an object if another DB2 holds retained locks on the object or has long-running noncommitting applications that use the object. You can use the -DISPLAY GROUP command to determine whether a member's status is ″FAILED.″ You can use the -DISPLAY DATABASE command with the LOCKS option to determine if locks are held.

RECOVER-pending and REBUILD-pending status: You cannot reorganize an index if any partition of the index is in the RECOVER pending status or in the REBUILD-pending status. Similarly, you cannot reorganize a single index partition if it is in the RECOVER-pending status or in the REBUILD-pending status.

There is one RECOVER-pending restrictive state:

RECP The index space or partition is in a RECOVER-pending status. A single logical partition in RECP does not restrict access to other logical partitions not in RECP. RECP can be reset by recovering only the single logical partition.

There are three REBUILD-pending restrictive states:

RBDP REBUILD-pending status (RBDP) is set on a physical or logical index partition. The individual physical or logical partition is inaccessible and must be rebuilt using the REBUILD INDEX utility.

PSRBD

Page set REBUILD-pending (PSRBD) is set for nonpartitioning indexes. The entire index space is inaccessible and must be rebuilt using the REBUILD utility.

RBDP*

A REBUILD-pending status that is set only on logical partitions of nonpartitioning indexes. The entire index is inaccessible, but is made available again when the affected partitions are rebuilt using the REBUILD INDEX utility.

For information about resetting the REBUILD-pending and RECOVER-pending states, see Table 134 on page 681 and Table 133 on page 680.

CHECK-pending status: You cannot reorganize an index when the data is in the CHECK-pending status. See "Chapter 6. CHECK DATA" on page 53 for more information about resetting the CHECK-pending status.

Data sets used by REORG INDEX

Table 67 on page 361 describes the data sets used by REORG. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 58. Data sets used by REORG INDEX

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
Work data set	One temporary data set for unload output and build input. The DD or TEMPLATE name is specified with the WORKDDN option of the utility control statement. The default DD name for this data set is SYSUT1.	Yes

The following objects are named in the utility control statement and do not require DD cards in the JCL:

Index The name of the index space to be reorganized. It is named in the control statement and is accessed through the DB2 catalog.

Calculating the size of the work data sets: When reorganizing an index space, you need a non-DB2 sequential work data set. That data set is identified by the DD statement named in the WORKDDN option. During the UNLOAD phase, the index keys and the data pointers are unloaded to the work data set. This data set is used to build the index. It is required only during the execution of REORG.

To calculate the approximate size (in bytes) of the WORKDDN data set SYSUT1, follow these steps:

1. Calculate the number of keys:

   ```
   number of keys = #tablerows
   ```

   where:

   #tablerows Number of records in the table.
2. Multiply the number of keys by the key length plus 9.

Creating the control statement

See "Syntax and options of the control statement" on page 304 for REORG syntax and option descriptions. See "Sample control statements" on page 325 for examples of REORG usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

"Determining when an object should be reorganized" on page 317

"Determining when an index requires reorganization"
"Specifying access with SHRLEVEL" on page 318
"Unloading without reloading" on page 319
"Considerations for fallback recovery" on page 319
"Changing data set definitions" on page 319
"Temporarily interrupting REORG" on page 320
"Improving performance" on page 320
"Improving performance with LOAD or REORG PREFORMAT" on page 197

Determining when an object should be reorganized

Product-sensitive Programming Interface

You can determine when to run REORG for indexes by using the LEAFDISTLIMIT catalog query option. If you specify the REPORTONLY option, REORG will produce a report detailing if a REORG is recommended; a REORG is not performed.

When you specify the catalog query options along with the REPORTONLY option, REORG produces a report with one of the following return codes:

1 No limit met; no REORG performed or recommended.
2 REORG performed or recommended.

Alternatively, information from the SYSINDEXPART catalog table can tell you which indexes qualify for reorganization.

End of Product-sensitive Programming Interface

Determining when an index requires reorganization

Product-sensitive Programming Interface

Use this query to identify user created indexes and DB2 catalog indexes that you should consider reorganizing with the REORG utility:

```
SELECT IXNAME, IXCREATOR
 FROM SYSIBM.SYSINDEXPART
   WHERE LEAFDIST > 200;
```

End of Product-sensitive Programming Interface

Be aware that using a LEAFDIST value of more than 200 as an indicator of a disorganized index is merely a rough rule of thumb for general cases. It is not absolute. There are cases where 200 is an acceptable value for LEAFDIST. For example, with FREEPAGE 0 and index page splitting, the LEAFDIST value can climb sharply. In this case, a LEAFDIST value higher than 200 can be acceptable.

Product-sensitive Programming Interface

After you run RUNSTATS, the following SQL statement provides the average distance (multiplied by 100) between successive leaf pages during sequential access of the ZZZ index.

```
SELECT LEAFDIST
  FROM SYSIBM.SYSINDEXPART
  WHERE IXCREATOR = 'index_creator_name'
  AND IXNAME = 'index_name';
```

End of Product-sensitive Programming Interface

If LEAFDIST increases over time, this probably indicates that the index should be reorganized. The optimal value of the LEAFDIST catalog column is zero. However, immediately after you run the REORG and RUNSTATS utilities, LEAFDIST might be greater than zero, due to empty pages for FREEPAGE and non-leaf pages.

Specifying access with SHRLEVEL

For reorganizing an index, or a partition of a index, the SHRLEVEL option lets you choose the level of access you have to your data during reorganization:

- REORG with SHRLEVEL NONE, the default, reloads the reorganized data into the original area being reorganized. Applications have read-only access during unloading and no access during reloading. SHRLEVEL NONE is the only access level that resets REORG-pending status.
- REORG with SHRLEVEL REFERENCE reloads the reorganized data into a new (shadow) copy of the area being reorganized. Near the end of reorganization, DB2 switches applications' future access from the original to the shadow copy. For SHRLEVEL REFERENCE, applications have read-only access during unloading and reloading, and a brief period of no access during switching.
- REORG with SHRLEVEL CHANGE reloads the reorganized data into a shadow copy of the area being reorganized. Applications can read and write the original area, and DB2 records the writing in the log. DB2 then reads the log and applies it to the shadow copy to bring the shadow copy up to date. This step executes iteratively, with each iteration processing a sequence of log records. Near the end of reorganization, DB2 switches applications' future access from the original to the shadow copy. Applications have read/write access during unloading and reloading, a brief period of read-only access during the last iteration of log processing, and a brief period of no access during switching.

Log processing with SHRLEVEL CHANGE: When you specify SHRLEVEL CHANGE, DB2 processes the log to update the shadow copy. This step executes iteratively. The first iteration processes the log records that accumulated during the previous iteration. The iterations continue until one of these conditions is met:

- DB2 estimates that the time to perform the log processing in the next iteration will be less than or equal to the time specified by MAXRO. If this condition is met, the next iteration will be the last.
- DB2 estimates that the SWITCH phase will not start by the deadline specified by DEADLINE. If this condition is met, DB2 terminates reorganization.
- The number of log records that the next iteration will process is not sufficiently lower than the number of log records processed in the previous iteration. If this condition is met but the first two conditions are not, DB2 sends message DSNU377I to the console. DB2 continues log processing for the length of time specified by DELAY and then performs the action specified by LONGLOG.

Operator actions: LONGLOG specifies the action that DB2 performs if log processing is not catching up. See "Option descriptions" on page 306 for a description of the LONGLOG options. If no action is taken after message DSNU377I is sent to the console, the LONGLOG option automatically goes into effect. Some examples of possible actions you may take:

- Execute the START DATABASE(db) SPACENAM(ts) ... ACCESS(RO) command and the QUIESCE utility to drain the write claim class. DB2 performs the last iteration, if MAXRO is not DEFER. After the QUIESCE, you should also execute the ALTER UTILITY command, even if you do not change any REORG parameters.

- Execute the START DATABASE(db) SPACENAM(ts) ... ACCESS(RO) command and the QUIESCE utility to drain the write claim class. Then, after reorganization has made some progress, execute the START DATABASE(db) SPACENAM(ts) ... ACCESS(RW) command. This increases the likelihood that log processing will catch up. After the QUIESCE, you should also execute the ALTER UTILITY command, even if you do not change any REORG parameters.
- Execute the ALTER UTILITY command to change the value of MAXRO. Changing it to a huge positive value, such as 9999999, causes the next iteration to be the last iteration.
- Execute the ALTER UTILITY command to change the value of LONGLOG.
- Execute the TERM UTILITY command to terminate reorganization.
- Adjust the amount of buffer space allocated to reorganization and to applications. This can increase the likelihood that log processing will catch up. After adjusting the space, you should also execute the ALTER UTILITY command, even if you do not change any REORG parameters.
- Adjust the scheduling priorities of reorganization and applications. This can increase the likelihood that log processing will catch up. After adjusting the priorities, you should also execute the ALTER UTILITY command, even if you do not change any REORG parameters.

DB2 does not take the action specified in the LONGLOG phrase if any one of these events occurs before the delay expires:

- An ALTER UTILITY command is issued.
- A TERM UTILITY command is issued.
- DB2 estimates that the time to perform the next iteration will be less than or equal to the time specified in the MAXRO phrase.
- REORG terminates for any reason (including the deadline).

Unloading without reloading

REORG can unload data without continuing and without leaving a SYSIBM.SYSUTIL record after the job ends.

Considerations for fallback recovery

Successful REORG INDEX processing inserts an SYSIBM.SYSCOPY row with ICTYPE='W' for an index that was defined with COPY YES. REORG also places a reorganized index in informational COPY-pending status. You should take a full image copy of the index after the REORG job completes to create a valid point of recovery.

Changing data set definitions

If the index space is defined by storage groups, space allocation is handled by DB2 and data set definitions cannot be altered during the reorganization process. DB2 deletes and redefines the necessary data sets to reorganize the object.

For REORG with SHRLEVEL REFERENCE or CHANGE, you can use the ALTER STOGROUP command to change the characteristics of a DB2-managed data set. The user effectively changes the characteristics of a user-managed data set by specifying the desired new characteristics when creating the shadow data set; see page 314 for more information about user-managed data sets. In particular, placing the original and shadow data sets on different DASD volumes might reduce contention and thus improve the performance of REORG and the performance of applications during REORG execution.

Temporarily interrupting REORG

You can temporarily pause REORG. If you specify UNLOAD PAUSE, REORG pauses after unloading the index space into the work data set. The job completes with return code 4. You can restart REORG using the phase restart or current restart. The REORG statement must not be altered.

The SYSIBM.SYSUTIL record for the REORG utility remains in ″stopped″ status until REORG is restarted or terminated.

While REORG is interrupted by PAUSE, you may re-define the table space attributes for user defined table spaces. PAUSE is not required for STOGROUP defined table spaces. Attribute changes are done automatically by a REORG following an ALTER INDEX.

Improving performance

To improve REORG performance:

- Run REORG concurrently on separate partitions of a partitioned index space. The processor time it takes to run REORG INDEX on partitions of a partitioned index is roughly the same as it would take to run a single REORG index job. The elapsed time is a fraction of what it would take to run a single REORG job on the entire index.

When to use SHRLEVEL CHANGE: Schedule REORG with SHRLEVEL CHANGE when the rate of writing is low and transactions are short. Avoid scheduling REORG with SHRLEVEL CHANGE when low-tolerance applications are executing. ***When to use DRAIN_WAIT:*** The DRAIN_WAIT option is provided to give greater control over the time online REORG will wait for drains. Also because the DRAIN_WAIT is the aggregate time that online REORG will wait to perform a drain on a table space and associated indexes, the length of drains will be more predictable than each part and index having it's own individual waiting time limit.

By specifying a short delay time (less than the system timeout value, IRLMRWT), the impact on applications can be reduced by reducing timeouts. The RETRY option could be used to give the online REORG further chances to complete successfully. If the user does not wish to use RETRY processing, DRAIN_WAIT can still be used to set a specific and more consistent limit on the length of drains.

RETRY allows an online REORG that is unable to drain the objects that it requires, to try again after a set period (RETRY_DELAY). During the RETRY_DELAY period, all the objects are available for READ/WRITE in the case of SHRLEVEL CHANGE. For SHRLEVEL REFERENCE, the objects remain with the access prior to the attempted drain (i.e. if the drain fails in the UNLOAD phase the object remains R/W, if the drain fails in the SWITCH phase, then objects remain in READ ONLY). Since application SQL can queue up behind any unsuccessful drain the online REORG has tried, it is recommended to have a reasonable delay before retrying, to allow this work to complete: the default is 5 minutes.

When DRAIN WRITERS (the default) is used with SHRLEVEL CHANGE and RETRY, multiple read only log iterations can occur. Since generally, online REORG may have to do more work when RETRY is specified and this may result in multiple or extended periods of restricted access. It is essential that applications running alongside online REORG, perform frequent commits. During the interval between retries, the utility is still active and consequently other utility activity against the table space and/or indexes are restricted.

When doing a table space REORG with RETRY and SHRLEVEL CHANGE both specified, the size of the COPY taken by REORG may increase.

Recommendation: Run online REORG during lighter periods of activity on the table space or index.

Terminating or restarting REORG INDEX

If you terminate REORG with the TERM UTILITY command during the UNLOAD phase, objects have not yet been changed and the job can be rerun.

If you terminate REORG with the TERM UTILITY command during the BUILD phase, the behavior depends on the SHRLEVEL option:

- For SHRLEVEL NONE, the index is left in RECOVER-pending status. After you recover the index, rerun the REORG job.
- For SHRLEVEL REFERENCE or CHANGE, the index keys are reloaded into a shadow index, so the original index has not been affected by REORG. You can rerun the job.

If you terminate REORG with the TERM UTILITY command during the LOG phase, the index keys are reloaded into a shadow index, so the original index has not been affected by REORG. You can rerun the job.

If you terminate REORG with the TERM UTILITY command during the SWITCH phase, all data sets that were renamed to their shadow counterparts are renamed back, so the objects are left in their original state. You can rerun the job. If there is a problem in renaming to the original data sets, the objects are left in RECOVER-pending status. You must recover the index.

The REORG-pending status is not reset until the UTILTERM execution phase. If the REORG INDEX utility abends or is terminated, the objects are left in RECOVER-pending status. See “Appendix C. Resetting an advisory or restrictive status” on page 677 for information about resetting either status.

Table 59. REORG INDEX phases and restrictive statuses

Phase	Effect on restrictive status
UNLOAD	No effect
BUILD	Sets REBUILD-pending (RBDP) status at the beginning of the BUILD phase, and resets RBDP at the end of the phase. SHRLEVEL NONE places an index that was defined with the COPY YES attribute in RECOVER pending (RECP) status.
LOG	No effect
SWITCH	Under certain conditions, if TERM UTILITY is issued, it must complete successfully or objects may be placed in RECP status or RBDP status. For SHRLEVEL REFERENCE or CHANGE, sets the RECP status if the index was defined with the COPY YES attribute at the beginning of the SWITCH phase, and resets RECP at the end of the phase. If the index was defined with COPY NO, sets the index in RBDP status at the beginning of the phase, and resets RBDP at the end of the phase.

Restarting REORG: Table 60 on page 322 provides information about restarting REORG INDEX.

If you restart REORG in the UTILINIT phase, it re-executes from the beginning of the phase. If REORG abends or system failure occurs while it is in the UTILTERM phase, you must restart the job with RESTART(PHASE).

For each phase of REORG and for each type of REORG INDEX (with SHRLEVEL NONE, with SHRLEVEL REFERENCE, and with SHRLEVEL CHANGE), the table indicates the types of restart allowed (CURRENT and PHASE). ″None″ indicates that no restart is allowed. A blank indicates that the phase does not occur. The ″Data Sets Required″ column lists the data sets that must exist to perform the specified type of restart in the specified phase.

Table 60. REORG INDEX utility restart information

Phase	Type for NONE	Type for REFERENCE	Type for CHANGE	Data Sets Required	Notes
UNLOAD	CURRENT PHASE	CURRENT PHASE	None None	SYSUT1	
BUILD	CURRENT PHASE	CURRENT PHASE	None None	SYSUT1 SYSUT1	1
LOG			None None		
SWITCH		CURRENT PHASE	CURRENT PHASE	originals and shadows originals and shadows	1

Notes:

1. You can restart the utility with either RESTART or RESTART(PHASE). However, because this phase does not take checkpoints, RESTART always re-executes from the beginning of the phase.

If you restart a REORG STATISTICS job using RESTART CURRENT, inline statistics collection will not occur. To update catalog statistics, run the RUNSTATS utility after the restarted job completes. Restarting a REORG STATISTICS job with RESTART(PHASE) is conditional after executing UNLOAD PAUSE. To determine if catalog table statistics will be updated using RESTART(PHASE), see Table 61.

Table 61. REORG INDEX utility phase restart using STATISTICS keyword

Phase	CURRENT	PHASE
UTILINIT	NO	YES
UNLOAD	NO	YES
BUILD	NO	YES

For instructions on restarting a utility job, see “Chapter 3. Invoking DB2 online utilities” on page 23.

Restarting REORG after an out of space condition: See “Restarting after the output data set is full” on page 47 for guidance in restarting REORG from the last commit point after receiving an out of space condition.

Concurrency and compatibility

Individual index partitions are treated as distinct target objects. Utilities operating on different partitions of the same index space are compatible.

REORG INDEX compatibility

Table 62 shows which claim classes REORG INDEX drains and any restrictive state the utility sets on the target object. The target is an index or index partition.

Table 63 shows which utilities can run concurrently with REORG INDEX on the same target object. The target object can be an index space or a partition. If compatibility depends on particular options of a utility, that is also shown.

Table 62. Claim classes of REORG INDEX operations. Use of claims and drains; restrictive states set on the target object.

Phase	REORG INDEX SHRLEVEL NONE	REORG INDEX SHRLEVEL REFERENCE	REORG INDEX SHRLEVEL CHANGE
UNLOAD	DW/UTRO	DW/UTRO	CR/UTRW
BUILD	DA/UTUT		
Last iteration of LOG	n/a	DA/UTUT[1]	DW/UTRO
SWITCH	n/a	DA/UTUT	DA/UTUT

Legend:
- CR: Claim the read claim class
- DA: Drain all claim classes, no concurrent SQL access
- DR: Drain the repeatable read class, no concurrent access for SQL repeatable readers
- DW: Drain the write claim class, concurrent access for SQL readers
- UTRO: Utility restrictive state, read only access allowed
- UTUT: Utility restrictive state, exclusive control
- Blank: Any claim, drain, or restrictive state for this object does not change in this phase.

Notes:

1. Applicable if you specified DRAIN ALL.

REORG INDEX does not set a utility restrictive state if the target object is an index on DSNDB01.SYSUTILX.

Table 63. REORG INDEX compatibility

Action	REORG INDEX SHRLEVEL NONE, REFERENCE, or CHANGE
CHECK DATA	No
CHECK INDEX	No
CHECK LOB	Yes
COPY INDEXSPACE	No
COPY TABLESPACE	Yes
DIAGNOSE	Yes
LOAD	No
MERGECOPY	Yes
MODIFY	Yes
QUIESCE	No
REBUILD INDEX	No
RECOVER INDEX	No
RECOVER INDEXSPACE	No
RECOVER TABLESPACE (no option)	Yes

Table 63. REORG INDEX compatibility (continued)

Action	REORG INDEX SHRLEVEL NONE, REFERENCE, or CHANGE
RECOVER TABLESPACE TOCOPY or TORBA	No
RECOVER TABLESPACE ERROR RANGE	Yes
REORG INDEX SHRLEVEL NONE, REFERENCE, or CHANGE	No
REORG TABLESPACE SHRLEVEL NONE UNLOAD CONTINUE or PAUSE, REORG SHRLEVEL REFERENCE, or REORG SHRLEVEL CHANGE	No
REORG TABLESPACE SHRLEVEL NONE UNLOAD ONLY or EXTERNAL without cluster index	Yes
REORG TABLESPACE SHRLEVEL NONE UNLOAD ONLY or EXTERNAL with cluster index	No
REPAIR LOCATE KEY	No
REPAIR LOCATE RID DUMP, VERIFY, or REPLACE	Yes
REPAIR LOCATE RID DELETE	No
REPAIR LOCATE TABLESPACE PAGE REPLACE	Yes
REPAIR LOCATE INDEX PAGE REPLACE	No
REPORT	Yes
RUNSTATS INDEX	No
RUNSTATS TABLESPACE	Yes
STOSPACE	Yes
UNLOAD	Yes

To run on SYSIBM.DSNLUX01 or SYSIBM.DSNLUX02, REORG INDEX must be the only utility in the job step and the only utility running in the DB2 subsystem.

Reviewing REORG INDEX output

The output from REORG INDEX consists of a reorganized index or index partition. Table 64 summarizes the effect of REORG.

Table 64. REORG INDEX summary

Specification	Results
REORG INDEX	Entire index (all parts if partitioned)
REORG INDEX PART *n*	Part *n* of partitioning index

When reorganizing an index, REORG leaves free pages and free space on each page in accordance with the current values of the FREEPAGE and PCTFREE parameters. (Those values can be set by the CREATE INDEX, or ALTER INDEX statements). REORG leaves one free page after reaching the FREEPAGE limit for each table in the index space.

Sample control statements

Example 1: Reorganizing an index. Reorganize index XMSGTXT1. Stop the utility after the index keys have been unloaded, but allow for subsequent restart.

```
REORG INDEX DSN8710.XMSGTXT1
  UNLOAD PAUSE
```

Example 2: REORG INDEX using STATISTICS. Reorganize the index XEMPL1, using the STATISTICS option to update the catalog table statistics for this index.

```
REORG INDEX DSN8710.XEMPL1
   SHRLEVEL REFERENCE STATISTICS
```

Example 3: Update statistics and history tables used for path selection. Reorganize index IUOE0801, using options to update catalog table statistics and history tables used for access path selection. Report the collected statistics and route the statistics to SYSPRINT.

```
REORG INDEX IUOE0801
        STATISTICS
        KEYCARD
        REPORT YES
        UPDATE ACCESSPATH
        HISTORY ACCESSPATH
```

Chapter 23. REORG TABLESPACE

The REORG TABLESPACE utility reorganizes a table space to improve access performance and reclaim fragmented space. In addition, the utility can reorganize a single partition or range of partitions of a partitioned table space. You can specify the degree of access to your data during reorganization, and collect inline statistics using the STATISTICS keyword. If you specify REORG TABLESPACE UNLOAD EXTERNAL, the data is unloaded in a format that is acceptable to the LOAD utility of any DB2 subsystem. You can also delete rows during the REORG by specifying the DISCARD option.

You can determine when to run REORG for non-LOB table spaces by using the OFFPOSLIMIT or INDREFLIMIT catalog query options. If you specify the REPORTONLY option, REORG will produce a report detailing if a REORG is recommended; a REORG is not performed.

Run the REORG TABLESPACE utility on a LOB table space to help increase the effectiveness of prefetch. For a LOB table space, REORG TABLESPACE performs these actions:

- Removes imbedded free space.
- Attempts to make LOB pages contiguous.

For a diagram of REORG TABLESPACE syntax and a description of available options, see “Syntax and options of the control statement” on page 329. For detailed guidance on running this utility, see “Instructions for running REORG TABLESPACE” on page 356.

Output: If the table space or partition has the COMPRESS YES attribute, then the data is compressed when reloaded. If you specify the KEEPDICTIONARY option of REORG, the current dictionary is used; otherwise a new dictionary is built.

REORG TABLESPACE can be executed on the table spaces in the DB2 catalog database (DSNDB06) and some table spaces in the directory database (DSNDB01). It cannot be executed on any table space in the DSNDB07 database.

Table 65. Summary of REORG TABLESPACE output

REORG Specified	Results
REORG TABLESPACE	Reorganizes all data, entire partitioning index, and all non-partitioning indexes.
REORG TABLESPACE PART *n*	Reorganizes data for PART *n*, PART *n* of the partitioning index, and index entries for PART *n* in all nonpartitioning indexes.
REORG TABLESPACE PART *n:m*	Reorganizes data for PART *n* through PART *m*, parts *n* through *m* of the partitioning index, and index entries for those parts in all nonpartitioning indexes.

Authorization required: To execute this utility on a user table space, the privilege set of the process must include one of the following:

- REORG privilege for the database
- DBADM or DBCTRL authority for the database
- SYSCTRL authority
- SYSADM authority.

© Copyright IBM Corp. 1983, 2001

To execute this utility on a table space in the catalog or directory, the privilege set of the process must include one of the following:
- REORG privilege for the DSNDB06 (catalog) database
- DBADM or DBCTRL authority for the DSNDB06 (catalog) database
- Installation SYSOPR authority
- SYSCTRL authority
- SYSADM or Installation SYSADM authority

If you use REORG TABLESPACE SHRLEVEL CHANGE, the privilege set must include DELETE, INSERT, SELECT, and UPDATE privileges on the mapping table (see page 357).

To run REORG TABLESPACE STATISTICS REPORT YES, the privilege set must include the SELECT privilege on the catalog tables.

An authority other than installation SYSADM or installation SYSOPR can receive message DSNT500I, “resource unavailable”, while trying to reorganize a table space in the catalog or directory. This can happen when the DSNDB06.SYSDBAUT or DSNDB06.SYSUSER catalog table space or one of the indexes is unavailable. If this problem occurs, run the REORG TABLESPACE utility again using an authorization ID with the installation SYSADM or installation SYSOPR authority.

Execution phases of REORG TABLESPACE: The REORG TABLESPACE utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup
UNLOAD	Unloads table space; sorts data if a clustering index exists and you specified either SORTDATA or SHRLEVEL CHANGE. If you specified NOSYSREC, passes rows in memory to the RELOAD phase, otherwise writes them to a sequential data set.
RELOAD	Reloads from the sequential data set into the table space; creates full image copies if you specify COPYDDN, RECOVERYDDN, SHRLEVEL REFERENCE, or SHRLEVEL CHANGE. If you specify the SORTKEYS option, a subtask sorts the index keys. Updates table and table space statistics.
SORT	Sorts index keys. If you specify the SORTKEYS option, the sorted keys are passed in memory to the BUILD phase.
BUILD	Builds indexes. Updates index statistics.
SORTBLD	If you specify a parallel index build using the SORTKEYS keyword, all activities that normally occur in both the SORT and BUILD phases occur in the SORTBLD phase instead.
LOG	Processes log iteratively; appends changed pages to the full image copies. Used only if you specify SHRLEVEL CHANGE.
SWITCH	Switches access to shadow copy of table space or partition. Used only if you specify SHRLEVEL REFERENCE or CHANGE.
BUILD2	Corrects nonpartitioning indexes if you specify REORG TABLESPACE PART SHRLEVEL REFERENCE or CHANGE.
UTILTERM	Cleanup

Execution phases of REORG TABLESPACE on a LOB table space: The REORG TABLESPACE utility operates in these three phases when you run it on a LOB table space:

Phase	Description
UTILINIT	Initialization and setup
REORGLOB	Rebuilds the LOB table space in place; no LOBs are unloaded or

reloaded. The LOB table space is set to RECOVER-pending status at the start of processing; this status is reset on REORGLOB completion. If the REORGLOB phase fails, the LOB table space remains in RECOVER-pending status.

UTILTERM Cleanup

You cannot restart REORG TABLESPACE on a LOB table space in the REORGLOB phase. Before executing REORG TABLESPACE on a LOB table space defined with LOG NO, you should take a full image copy to ensure recoverability.

If the LOB table space was defined with LOG NO, it is left in COPY pending status after REORG TABLESPACE completes processing.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

REORG TABLESPACE

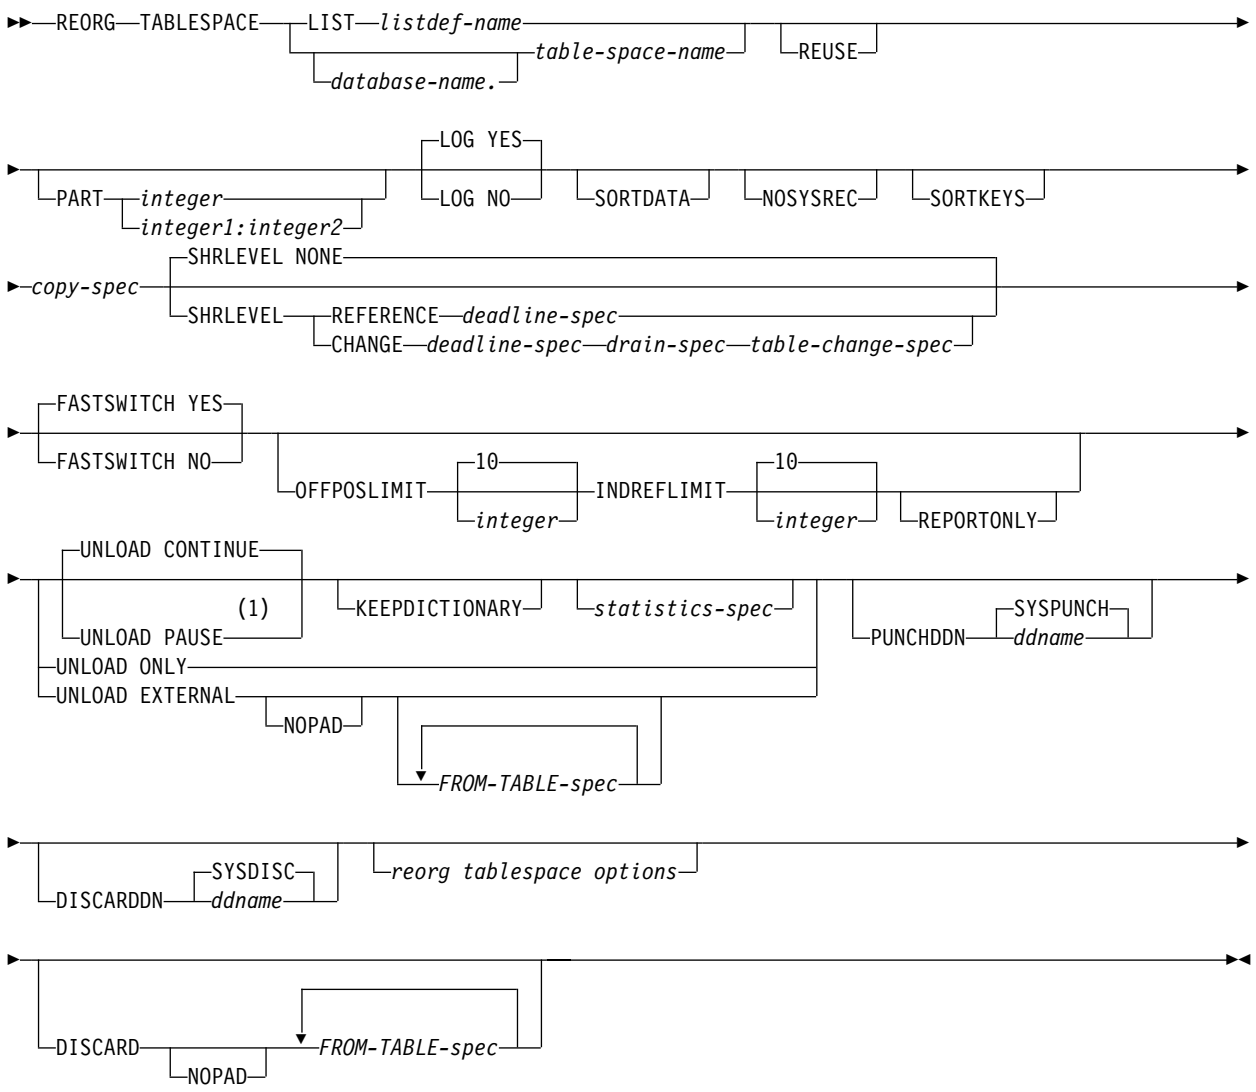

Notes:

1 You cannot use UNLOAD PAUSE with the LIST option.

copy-spec:

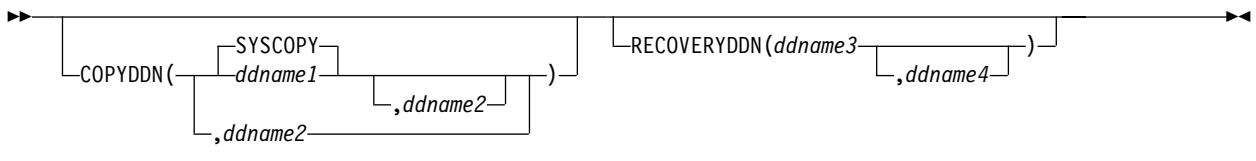

deadline-spec:

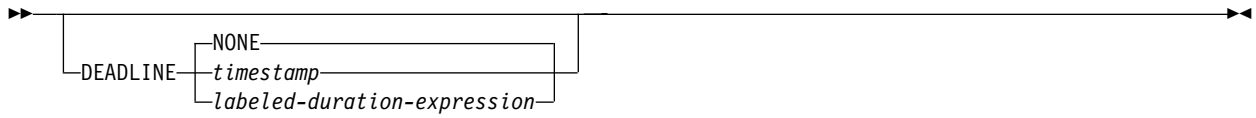

drain-spec:

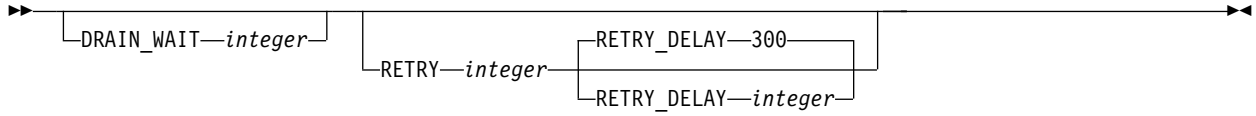

table-change-spec:

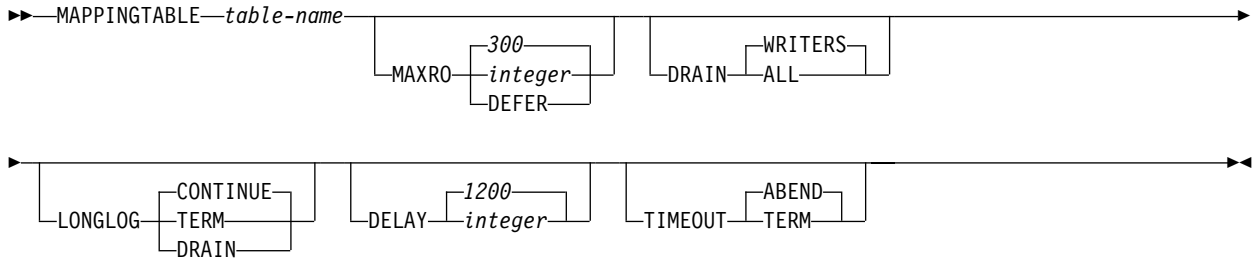

labeled-duration-expression:

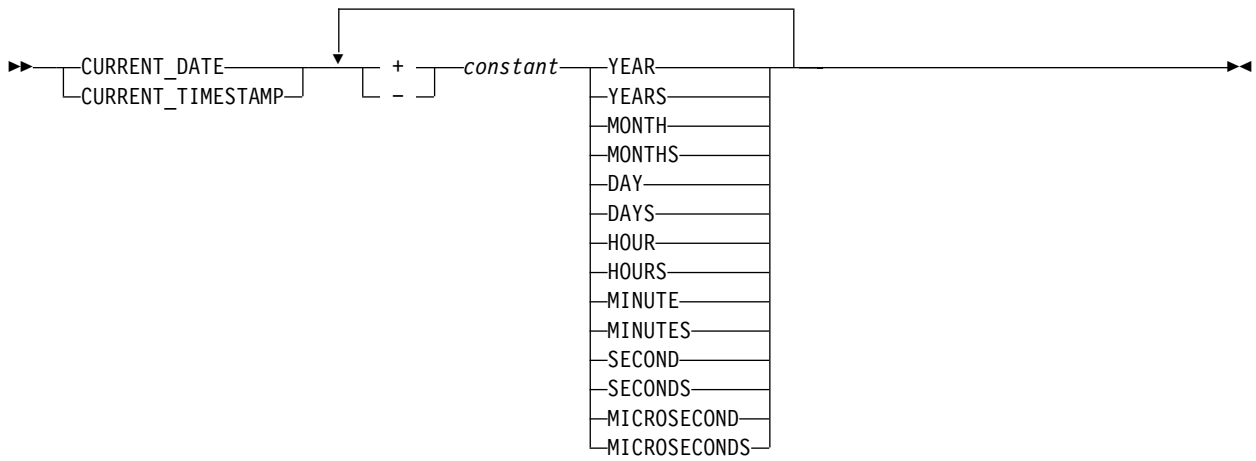

REORG TABLESPACE

statistics-spec:

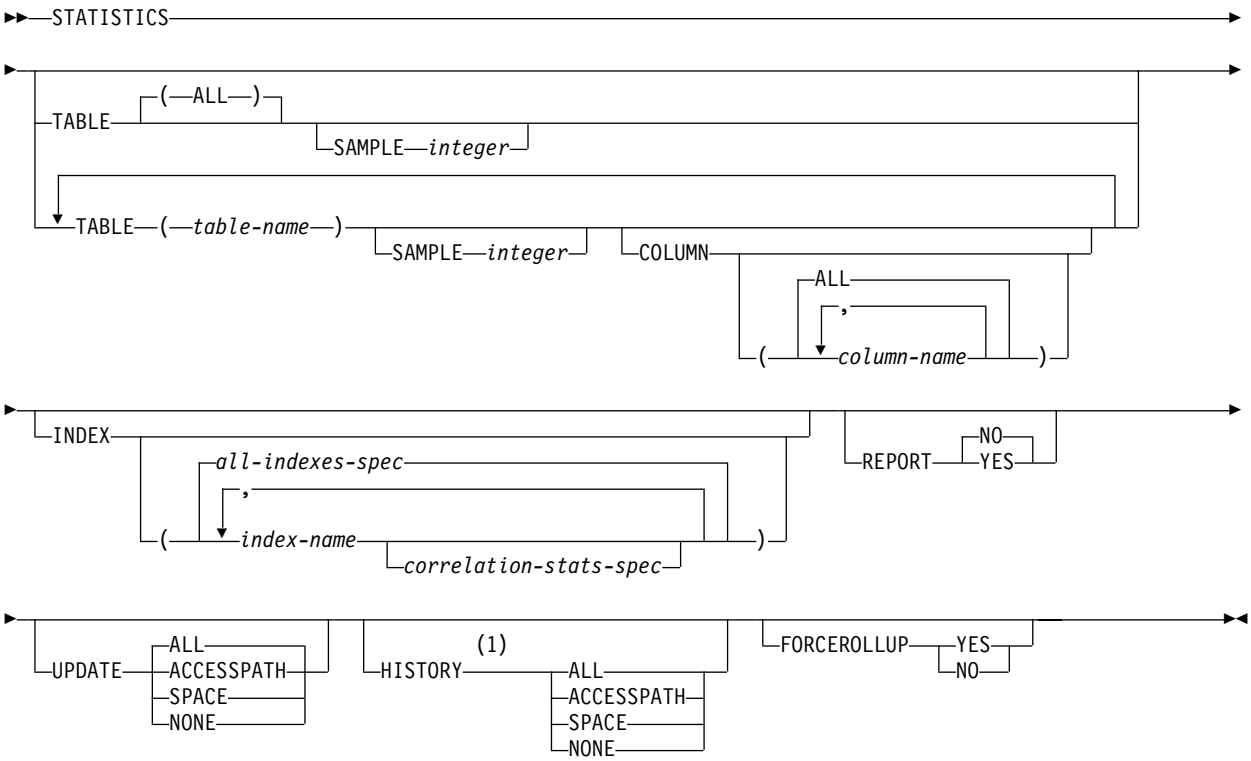

Notes:

1 What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

all-indexes-spec:

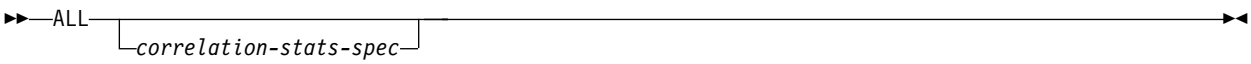

correlation-stats-spec:

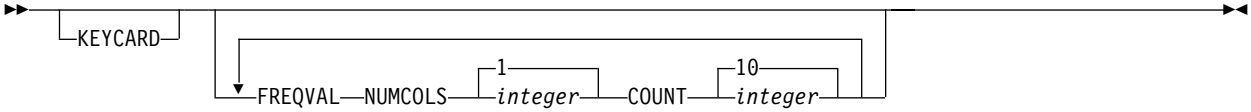

FROM-TABLE-spec:

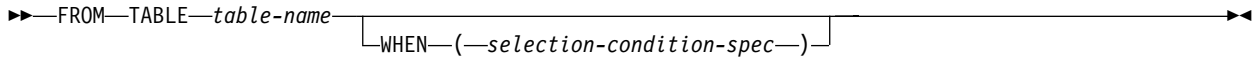

selection-condition-spec:

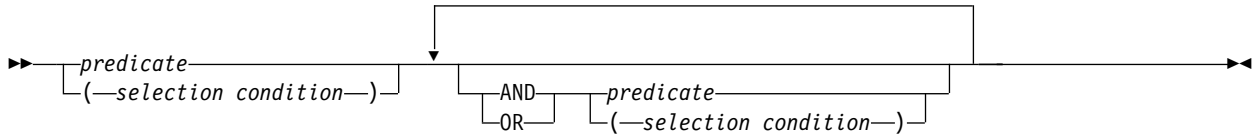

predicate:

basic predicate:

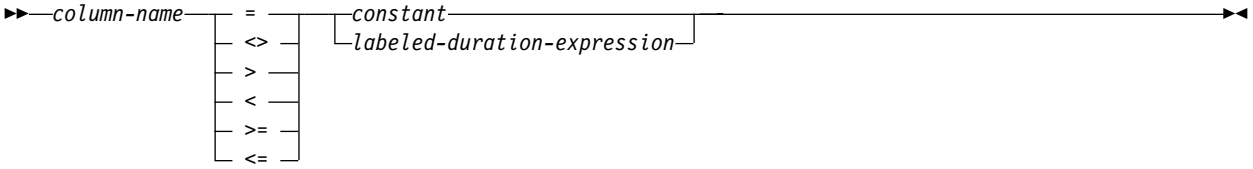

Note: The following forms of the comparison operators are also supported in basic and quantified predicates: !=, !<, and !>. In addition, in code pages 437, 819, and 850, the forms ¬=, ¬<, and ¬> are supported. All these product-specific forms of the comparison operators are intended only to support existing SQL statements that use these operators and are not recommended for use when writing new SQL statements.

A not sign (¬), or the character that must be used in its place in certain countries, can cause parsing errors in statements passed from one DBMS to another. The problem occurs if the statement undergoes character conversion with certain combinations of source and target CCSIDs. To avoid this problem, substitute an equivalent operator for any operator that includes a not sign. For example, substitute '< >' for '¬=', '<=' for '¬>', and '>=' for '¬<'.

BETWEEN predicate:

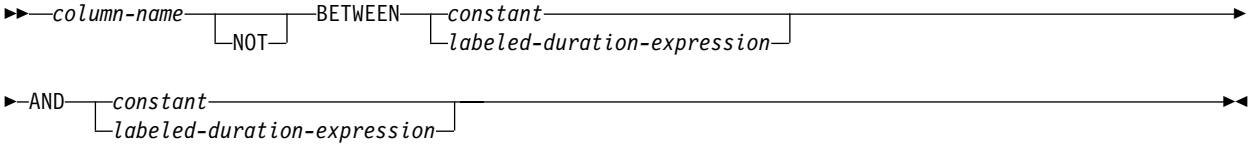

IN predicate:

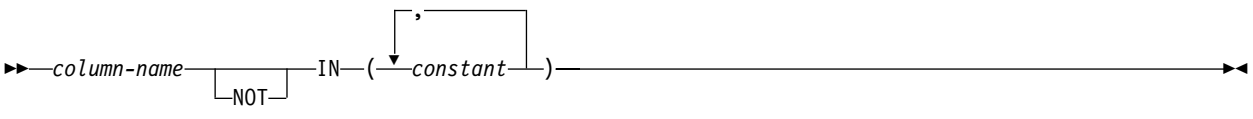

LIKE predicate:

NULL predicate:

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

TABLESPACE *database-name.table-space-name*
Specifies the table space (and, optionally, the database to which it belongs) that is to be reorganized.

If you reorganize a table space, its indexes are also reorganized.

database-name
Is the name of the database to which the table space belongs. The name cannot be DSNDB07.

The **default** is **DSNDB04**.

table-space-name
Is the name of the table space to be reorganized. The name cannot be SYSUTILX if the database name specified is DSNDB01.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each control statement of REORG TABLESPACE. The list must contain only table spaces.

Do not specify FROM TABLE, STATISTICS TABLE *table-name*, or STATISTICS INDEX *index-name* with REORG TABLESPACE LIST. If you want to collect inline statistics for a LIST of table spaces, specify STATISTICS TABLE (ALL). If you want to collect inline statistics for a LIST of indexes, specify STATISTICS INDEX (ALL). Do not specify PART with LIST.

Utility list manager will invoke REORG TABLESPACE once per item in the list.

For more information about LISTDEF specifications, see "Generating object lists using LISTDEF" on page 144.

REUSE
When used with SHRLEVEL NONE, specifies that REORG logically resets and reuses DB2-managed data sets without deleting and redefining them. If you do not specify REUSE and SHRLEVEL NONE, DB2 deletes and redefines DB2-managed data sets to reset them.

If a data set has multiple extents, the extents will not be released if you use the REUSE parameter.

REUSE does not apply if you also specified SHRLEVEL REFERENCE or CHANGE.

PART *integer*

PART *integer1*:*integer2*
Identifies a partition to be reorganized. You can reorganize a single partition of a partitioned table space, or a range of partitions within a partitioned table space. *integer* must be in the range from 1 to the number of partitions that are defined for the table space or partitioning index. The maximum is 254.

integer Designates a single partition.

integer1:integer2
Designates a range of existing table space partitions from *integer1* through *integer2*.

If you omit the PART keyword, the entire table space is reorganized.

If you specify the PART keyword for a LOB table space, DB2 issues an error message, and utility processing terminates with return code 8.

LOG
Specifies whether records are logged during the reload phase of REORG. If the records are not logged, the table space is recoverable only after an image copy has been taken. If you specify COPYDDN, RECOVERYDDN, SHRLEVEL REFERENCE, or SHRLEVEL CHANGE, an image copy is taken during REORG execution.

YES
Logs records during the reload phase. This option is not allowed for any table space in DSNDB01 or DSNDB06, or if the SHRLEVEL REFERENCE or CHANGE options are used.

If you specify SHRLEVEL NONE (explicitly or by default), the **default** is **YES**.

You must specify LOG YES (explicitly or by default) for a LOB table space. Logging will occur only if the LOB table space was defined with the LOG YES attribute. If the LOB table space was defined with the LOG NO attribute, the LOB table space will be left in COPY-pending status after the REORG.

NO
Does not log records. Puts the table space in COPY-pending status if either of these conditions is true:

- REORG is executed at the local site, and neither COPYDDN, SHRLEVEL REFERENCE, nor SHRLEVEL CHANGE are specified.
- REORG is executed at the remote site, and RECOVERYDDN is not specified.

SORTDATA
Specifies that the data is to be unloaded by table space scan, then sorted in clustering order. Records always are sorted by the table in order to retain the clustering of records of the same table.

This option is recommended to improve performance unless one of the following is true:

- The data is in perfect clustering order and the REORG utility is used to reclaim space from dropped tables.
- The data set is very large and there is not enough DASD available for sorting.
- The longest possible composite record size is greater than 32760.

SORTDATA is ignored for some the catalog and directory table spaces; see “Reorganizing the catalog and directory” on page 369.

NOSYSREC
Specifies that the output of sorting (if there is a clustering index and you specify SORTDATA) is the input to reloading, without using an unload data set. You can specify this option only if you specify REORG TABLESPACE, SORTDATA, SHRLEVEL REFERENCE, or SHRLEVEL NONE, and only if you do not specify UNLOAD PAUSE or UNLOAD ONLY. See “Omitting the output data set” on page 368 for additional information about using this option.

SORTKEYS
Specifies that index keys will be sorted and indexes will be built in parallel during the SORTBLD phase to improve performance. This option is recommended if more than one index needs to be created.

COPYDDN *ddname1,ddname2*
Specifies the DD statements for the primary (*ddname1*) and backup (*ddname2*) copy data sets for the image copy.

ddname is the DD name.

The **default** is **SYSCOPY** for the primary copy. A full image copy data set is created when REORG executes. In the row in the SYSIBM.SYSCOPY catalog table, the SHRLEVEL column is set to ″R,″ as it would be for the COPY SHRLEVEL REFERENCE. The table space is not left in COPY-pending status regardless of which LOG option you specify.

If you specify SHRLEVEL NONE (explicitly or by default) for REORG, and COPYDDN is not specified, then no image copy is created at the local site.

COPYDDN(SYSCOPY) is assumed, and a DD statement for SYSCOPY is required if:

- You specify REORG SHRLEVEL REFERENCE or CHANGE, and do not specify COPYDDN
- A partition is in REORG-pending (REORP) status

| The COPYDDN keyword specifies either a DDNAME or a TEMPLATE name
| specification from a previous TEMPLATE control statement. If utility processing
| detects that the specified name is both a DDNAME in the current job step and a
| TEMPLATE name, the utility uses DDNAME. For more information about
| TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

When using the inline COPY function, the FASTSWITCH methodology is used
in the SWITCH phase. The default for the FASTSWITCH parameter is YES. If
you do not specify FASTSWITCH NO, the fifth level qualifier in the data set
name alternates between I0001 and J0001.

RECOVERYDDN *(ddname3,ddname4)*
Specifies the DD statements for the primary *(ddname3)* and backup *(ddname4)* copy data sets for the image copy at the recovery site.

ddname is the DD name.

You cannot have duplicate image copy data sets. The same rules apply for RECOVERYDDN as for COPYDDN.

| The RECOVERYDDN keyword specifies either a DDNAME or a TEMPLATE
| name specification from a previous TEMPLATE control statement. If utility
| processing detects that the specified name is both a DDNAME in the current
| job step and a TEMPLATE name, the utility uses DDNAME. For more
| information about TEMPLATE specifications, see “Chapter 28. TEMPLATE” on page 463.

When using the inline COPY function, the FASTSWITCH methodology is used
in the SWITCH phase. The default for the FASTSWITCH parameter is YES. If
you do not specify FASTSWITCH NO, the fifth level qualifier in the data set
name alternates between I0001 and J0001.

SHRLEVEL

Specifies the method for performing the reorganization. The parameter following SHRLEVEL indicates the type of access allowed during the RELOAD phase of REORG.

For a LOB table space, you must specify SHRLEVEL NONE (explicitly or by default).

NONE Specifies that reorganization operates by unloading from the area being reorganized (while applications can read but cannot write to the area), reloading into that area (while applications have no access), and then allowing read/write access again.

The **default** is **NONE**.

If you specify NONE (explicitly or by default), the following parameters cannot be specified: MAPPINGTABLE, MAXRO, LONGLOG, DELAY, and DEADLINE. If you omit SORTDATA, or specify UNLOAD PAUSE or UNLOAD ONLY, then you cannot specify NOSYSREC.

REFERENCE

Specifies that reorganization operates by unloading from the area being reorganized (while applications can read but cannot write to the area), reloading into a shadow copy of that area (while applications can read but cannot write to the original copy), switching applications' future access from the original copy to the shadow copy by exchanging the names of the data sets, and then allowing read/write access again.

To determine which data sets are required when you execute REORG SHRLEVEL REFERENCE, see "Data sets used by REORG TABLESPACE" on page 361.

If you specify REFERENCE, you cannot specify the following parameters:

- LOG. Reorganization with REFERENCE always creates an image copy and always refrains from logging records during reloading.
- UNLOAD. Reorganization with REFERENCE always performs UNLOAD CONTINUE.
- MAPPINGTABLE, MAXRO, LONGLOG, and DELAY.

SHRLEVEL REFERENCE is not allowed for a LOB table space.

CHANGE

Specifies that reorganization operates by unloading from the area being reorganized (while applications can read and write to the area), reloading into a shadow copy of that area (while applications have read/write access to the original copy of the area), applying the log of the original copy to the shadow copy (while applications can read and usually write to the original copy), switching applications' future access from the original copy to the shadow copy by exchanging the names of the data sets, and then allowing read/write access again.

To determine which data sets are required when you execute REORG SHRLEVEL CHANGE, see "Data sets used by REORG TABLESPACE" on page 361.

If you specify CHANGE, you cannot specify the following parameters:

- LOG. Reorganization with CHANGE always creates an image copy and always refrains from logging records during reloading.

- SORTDATA, NOSYSREC, SORTKEYS. Reorganization with CHANGE always operates as if these parameters were specified.
- UNLOAD. Reorganization with CHANGE always performs UNLOAD CONTINUE.

SHRLEVEL CHANGE is not allowed for a LOB table space.

FASTSWITCH
Specifies whether SWITCH phase methodology is enabled.

YES
FASTSWITCH YES enables the SWITCH phase methodology. You have a choice between which methodology to use for a given reorganization. This option is not allowed for the catalog (DSNDB06) or directory (DSNDB01).

When using the inline COPY function, the FASTSWITCH methodology is used in the SWITCH phase. The default for the FASTSWITCH parameter is YES. If you do not specify FASTSWITCH NO, the fifth level qualifier in the data set name alternates between I0001 and J0001.

The **default** is **YES**.

NO
FASTSWITCH NO will cause the SWITCH phase to use IDCAMS RENAME.

DEADLINE
Specifies the deadline for the switch phase to finish. If DB2 estimates that the SWITCH phase will not complete by the deadline, DB2 issues the messages that the -DISPLAY UTILITY command would issue and then terminates reorganization.

NONE Specifies that there is no deadline by which the SWITCH phase of log processing must finish.

The **default** is **NONE**.

timestamp
timestamp specifies the deadline for the SWITCH phase of log processing to finish. This deadline must not have already occurred when REORG is executed.

labeled-duration-expression
Calculates the deadline for the SWITCH phase of log processing to finish. The calculation is either based on CURRENT TIMESTAMP or CURRENT DATE. This deadline must not have already occurred when REORG is executed.

For example, to ensure that the SWITCH phase is complete by 6:30 AM two days from now, use the following expression:

```
DEADLINE CURRENT DATE + 2 DAYS + 6 HOURS + 30 MINUTES
```

If REORG SHRLEVEL REFERENCE or SHRLEVEL CHANGE terminates due to a DEADLINE specification, DB2 issues message DSNU374I with reason code 2, but does not set a restrictive status.

DRAIN_WAIT *integer*
Specifies the number of seconds that the utility will wait for DML (inserts, updates, deletes and select / fetches) when draining. The time specified is the aggregate time for all partitions of the index to be reorganised. This overrides the values specified by IRLMRWT and UTIMOUT, for DML only. For other tasks such as COMMANDS, the IRLMRWT and UTIMOUT values are used. If the

keyword is omitted or a value of 0 is specified, then regular draining using the IRLMRWT and UTIMOUT values will be used. Acceptable values can be from 0 to 1800.

RETRY *integer*
Specifies the maximum number of retries that can be attempted. Values can be from 0 to 255. If the keyword is omitted, then no retries will be attempted.

Note: Specifying RETRY can lead to increased processing costs and can result in multiple or extended periods of read-only access. When RETRY and SHRLEVEL CHANGE are both specified, the size of the COPY taken by REORG may increase.

RETRY_DELAY *integer*
Specifies the minimum duration in seconds between retries. Values can be from 1 to 1800.

The **default** is **300** seconds. The value must be an integer.

MAPPINGTABLE *table-name*
Specifies the name of the mapping table that REORG TABLESPACE uses to map between the RIDs of data records in the original copy of the area and the corresponding RIDs in the shadow copy. This parameter is required if you specify REORG TABLESPACE SHRLEVEL CHANGE, and you must create a mapping table and an index for it before running REORG TABLESPACE. The table must have the columns and the index that appear in the SQL statements described on page 357.

MAXRO *integer*
Specifies the maximum amount of time for the last iteration of log processing. During that iteration, applications have read-only access.

The actual execution time of the last iteration might exceed the value specified with MAXRO.

The ALTER UTILITY command can change the value of MAXRO.

The **default** is **300** seconds. The value must be an integer.

integer
integer is the number of seconds. Specifying a small positive value reduces the length of the period of read-only access, but it might increase the elapsed time for REORG to complete. If you specify a huge positive value, the second iteration of log processing is probably the last iteration.

DEFER
Specifies that the iterations of log processing with read/write access can continue indefinitely. REORG never begins the final iteration with read-only access, unless you change the MAXRO value with ALTER UTILITY.

If you specify DEFER, you should also specify LONGLOG CONTINUE.

If you specify DEFER, and DB2 determines that the actual time for an iteration and the estimated time for the next iteration are both less than 5 seconds, DB2 adds a 5 second pause to the next iteration. This pause reduces consumption of processor time. The first time this situation occurs for a given execution of REORG, DB2 sends message DSNU362I to the console. The message states that the number of log records that must be processed is small and that the pause will occur. The message also suggests that this would be an appropriate time to

execute ALTER UTILITY to change the MAXRO value and thus cause REORG to finish. DB2 adds the pause whenever the situation occurs; however, DB2 sends the message only if 30 minutes have elapsed since the last message was sent for a given execution of REORG.

DRAIN
Specifies drain behavior at the end of the log phase after the MAXRO threshold is reached and when the last iteration of the log is to be applied.

WRITERS
Specifies the current default action, in which DB2 drains just writers during the log phase after the MAXRO threshold is reached and subsequently issues DRAIN ALL on entering the switch phase.

ALL Specifies that DB2 drain all readers and writers during the log phase, after the MAXRO threshold is reached.

Consider specifying DRAIN ALL if the following conditions are both true:

- There is a lot of SQL update activity during the log phase.
- The default behavior results in a large number of -911 SQL error messages.

LONGLOG
Specifies the action that DB2 performs, after sending a message to the console, if the number of records that the next iteration of log process will process is not sufficiently lower than the number that the previous iterations processed. This situation means that reorganization's reading of the log is not catching up to applications' writing of the log quickly enough.

CONTINUE
Specifies that until the time on the JOB statement expires, DB2 continues performing reorganization, including iterations of log processing, if the estimated time to perform an iteration exceeds the time specified with MAXRO.

A value of DEFER for MAXRO and a value of CONTINUE for LONGLOG together mean that REORG continues allowing access to the original copy of the area being reorganized and does not switch to the shadow copy. The user can execute the -ALTER UTILITY command with a large value for MAXRO when the switching is desired.

The **default** is **CONTINUE**.

TERM Specifies that DB2 will terminate reorganization after the delay specified by the DELAY parameter.

DRAIN
Specifies that DB2 drains the write claim class after the delay specified by the DELAY parameter. This action forces the final iteration of log processing to occur.

DELAY *integer*
Specifies the minimum interval between the time that REORG sends the LONGLOG message to the console and the time REORG that performs the action specified by the LONGLOG parameter.

integer is the number of seconds. The value must be an integer. The **default** is **1200**.

TIMEOUT
Specifies the action to be taken if the REORG utility gets a time out condition

while trying to drain objects in either the LOG or SWITCH phases.This option
only applies if the SHRLEVEL CHANGE or the SHRLEVEL REFERENCE
option is used.

ABEND
If a time out condition occurs, DB2 leaves the objects in a UTRO or UTUT state.

TERM
If you specify the TERM option and a time out condition occurs, then DB2:

1. Issues an implicit TERM UTILITY command, causing the utility to end with a return code 8
2. Issues the DSNU590I and DSNU170I messages
3. Leaves the objects in a RW state.

OFFPOSLIMIT
The specified value is compared to the result of the calculation (NEAROFFPOSF + FAROFFPOSF) × 100 / CARDF for the specified partitions in SYSIBM.SYSINDEXPART for the explicit clustering indexes for every table in the specified table space. Alternatively, the values in SYSINDEXPART are checked for a single non-partitioned table space, or for each partition if you specified an entire partitioned table space as the target object. If at least one calculated value exceeds the OFFPOSLIMIT value, REORG is performed or recommended. This option is valid for non-LOB table spaces only.

The **default** value is **10**.

INDREFLIMIT
The specified value is compared to the result of the calculation (NEARINDREF + FARINDREF) × 100 / CARDF for the specified partitions in SYSIBM.SYSTABLEPART for the specified table space. Alternatively, the values in SYSTABLEPART are checked for a single non-partitioned table space, or for each partition if you specified an entire partitioned table space as the target object. If at least one calculated value exceeds the INDREFLIMIT value, REORG is performed or recommended. This option is valid for non-LOB table spaces only.

The **default** value is **10**.

REPORTONLY
If you specify this option, the REORG is only recommended, not performed. REORG produces a report with one of the following return codes:

1 No limit met; no REORG performed or recommended.
2 REORG performed or recommended.

UNLOAD
Specifies whether the utility job is to continue processing or end after the data is unloaded. Unless you specify UNLOAD EXTERNAL, data can be reloaded only into the same table and table space (as defined in the DB2 catalog) on the same subsystem. (This does not preclude VSAM redefinition during UNLOAD PAUSE.)

You must specify UNLOAD ONLY for the data set to be in a format that is compatible with the FORMAT UNLOAD option of LOAD. However, with LOAD you can load the data only into the same object from which it is unloaded. This option is valid for non-LOB table spaces only.

You must specify UNLOAD EXTERNAL for the data set to be in a format that is usable by LOAD without the FORMAT UNLOAD option. With UNLOAD EXTERNAL, you can load the data into any table with compatible columns in any table space on any subsystem.

CONTINUE
Specifies that, after the data has been unloaded, the utility continues processing. An edit routine may be called to decode a previously encoded data row if an index key requires extraction from that row.

If you specify DISCARD, rows are decompressed and edit routine decoded. If you also specify DISCARD to a file, rows will be field procedure decoded, and SMALLINT, INTEGER, FLOAT, DECIMAL, DATE, TIME, and TIMESTAMP columns will be converted to external format. Otherwise, edit routines or field procedures are bypassed on both the UNLOAD and RELOAD phases for table spaces. Validation procedures are not invoked during either phase.

The **default** is **CONTINUE**.

PAUSE
Specifies that after the data has been unloaded, processing ends. The utility stops and the RELOAD status is stored in SYSIBM.SYSUTIL so that processing can be restarted with RELOAD RESTART(PHASE).

This option is useful if you want to redefine data sets during reorganization. For example, with a user defined data set, you can:
- Run REORG with the UNLOAD PAUSE option
- Redefine the data set using access method services
- Restart REORG by resubmitting the previous job and specifying RESTART(PHASE).

However, you cannot use UNLOAD PAUSE if you specify the LIST option.

If zero records are unloaded during an UNLOAD PAUSE, when REORG is restarted, the RELOAD and BUILD phases are bypassed.

If you specify DISCARD, rows are decompressed and edit routine decoded. If you also specify DISCARD to a file, rows will be field procedure decoded, and SMALLINT, INTEGER, FLOAT, DECIMAL, DATE, TIME, and TIMESTAMP columns will be converted to external format. Otherwise, edit routines or field procedures are bypassed on both the UNLOAD and RELOAD phases for table spaces. Validation procedures are not invoked during either phase.

ONLY Specifies that after the data has been unloaded, the utility job ends and the status in SYSIBM.SYSUTIL corresponding to this utility ID is removed.

If you specify UNLOAD ONLY with REORG TABLESPACE, any edit routine or field procedure is executed during record retrieval in the unload phase.

This option is not allowed for any table space in DSNDB01 or DSNDB06.

The DISCARD and WHEN options are not allowed with UNLOAD ONLY.

EXTERNAL

Specifies that after the data has been unloaded, the utility job is to end and the status in SYSIBM.SYSUTIL corresponding to this utility ID is removed.

The UNLOAD utility has more functions. If you specify UNLOAD EXTERNAL with REORG TABLESPACE, rows are decompressed, edit routines decoded, field procedures are decoded, and SMALLINT, INTEGER, FLOAT, DECIMAL, DATE, TIME, and TIMESTAMP columns are converted to external format. Validation procedures are not invoked.

This option is not allowed for any table space in DSNDB01 or DSNDB06.

The DISCARD option is not allowed with UNLOAD EXTERNAL.

KEEPDICTIONARY

Prevents REORG TABLESPACE from building a new compression dictionary when unloading the rows. The efficiency of REORG increases with the KEEPDICTIONARY option for the following reasons:

- The processing cost of building the compression dictionary is eliminated.
- Existing compressed rows do not have to be compressed again.
- Existing compressed rows do not have to be expanded, unless indexes require it or SORTDATA is used.

Possible reasons for omitting KEEPDICTIONARY are:

- If the data has changed significantly since the last dictionary was built, rebuilding the dictionary might save a significant amount of space.
- If the current dictionary was built using the LOAD utility, building it using REORG might produce a better compression dictionary.

For more information about specifying or omitting the KEEPDICTIONARY option, see “Compressing data” on page 192.

KEEPDICTIONARY is valid only if a compression dictionary exists and the table space or partition being reorganized has the COMPRESS YES attribute. If a dictionary does not exist, one is built, a warning message is issued, and all the records are compressed.

Messages DSNU234I and DSNU244I, which show compression statistics, are not issued when you specify REORG UNLOAD CONTINUE KEEPDICTIONARY or REORG UNLOAD PAUSE KEEPDICTIONARY.

REORG ignores the KEEPDICTIONARY option if a partition that is being reorganized is in REORG pending status.

For information regarding ESA data compression, see Part 5 (Volume 2) of *DB2 Administration Guide*.

STATISTICS

Specifies the gathering of statistics for the table space or associated index, or both; the statistics are reported or stored in the DB2 catalog.

If you specify a table space partition or a range of partitions along with the STATISTICS keyword, DB2 collects statistics only for the specified table space partitions. This option is valid for non-LOB table spaces only.

You cannot collect inline statistics for indexes on specific catalog and directory tables. See Limitations for reorganizing the catalog and directory on page 370 for the list of unsupported catalog and directory tables.

TABLE
Specifies the table for which column information is to be gathered. All tables must belong to the table space specified in the TABLESPACE option.

Do not specify STATISTICS TABLE *table-name* with REORG TABLESPACE LIST. Instead, specify STATISTICS TABLE (ALL).

(ALL)
Specifies that information is to be gathered for all columns of all tables in the table space.

SAMPLE *integer*
Indicates the percentage of rows to sample when collecting non-indexed column statistics. Any value from 1 through 100 can be specified. The **default** is **25**. The SAMPLE option is not allowed for LOB table spaces.

(*table-name***)**
Specifies the tables for which column information is to be gathered. The parentheses are required. If you omit the qualifier, the user identifier for the utility job is used.

If you specify more than one table, you must repeat the TABLE option.

COLUMN
Specifies columns for which column information is to be gathered.

You can only specify this option if you specify a particular tables for which statistics are to be gathered (TABLE (*table-name*)). If you specify particular tables and do not specify the COLUMN option, the default, COLUMN(ALL), is used. If you do not specify a particular table when using the TABLE option, you cannot specify the COLUMN option; however, COLUMN(ALL) is assumed.

(ALL)
Specifies that statistics are to be gathered for all columns in the table.

(column-name, ...)
Specifies the columns for which statistics are to be gathered. The parentheses are required.

You can specify a list of column names; the maximum is 10. If you specify more than one column, separate each name with a comma.

INDEX
Specifies indexes for which information is to be gathered. Column information is gathered for the first column of the index. All the indexes must be associated with the *same* table space, which must be the table space specified in the TABLESPACE option.

Do not specify STATISTICS INDEX *index-name* with REORG TABLESPACE LIST. Instead, specify STATISTICS INDEX (ALL).

(ALL) Specifies that the column information is to be gathered for all indexes defined on tables contained in the table space. The parentheses are required.

(*index-name***)**
Specifies the indexes for which information is to be gathered. The parentheses are required.

REPORT
Determines if a set of messages is generated to report the collected statistics.

NO
Indicates that the set of messages is not output to SYSPRINT.

The **default** is **NO**.

YES
Indicates that the set of messages is output to SYSPRINT. The messages generated are dependent on the combination of keywords (such as TABLESPACE, INDEX, TABLE, and COLUMN) specified with the RUNSTATS utility. However, these messages are *not* dependent on the specification of the UPDATE option. REPORT YES always generates a report of SPACE and ACCESSPATH statistics.

UPDATE
Tells whether the collected statistics are inserted into the catalog tables. UPDATE also allows you to select statistics used for access path selection or statistics used by database administrators. What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

ALL Indicates that all collected statistics will be updated in the catalog.

The **default** is **ALL**.

ACCESSPATH
Indicates that only the catalog table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only the catalog table columns that provide statistics to help the database administrator assess the status of a particular table space or index are updated.

NONE Indicates that no catalog tables are updated with the collected statistics. This option is only valid when REPORT YES is specified.

HISTORY
Records all catalog table inserts or updates to the catalog history tables. What you specify under the UPDATE option affects what options you can specify under the HISTORY option. The table below lists the HISTORY options allowed for each UPDATE option specified.

The default will be supplied by the value specified in STATISTICS HISTORY on panel DSNTIPO.

ALL Indicates that all collected statistics will be updated in the catalog history tables.

ACCESSPATH
Indicates that only the catalog history table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only SPACE-related catalog statistics are updated in catalog history tables.

NONE Indicates that no catalog history tables are updated with the collected statistics.

Table 66. Allowable HISTORY options based on selected UPDATE option

UPDATE option	HISTORY options
ALL	ALL, ACCESSPATH, SPACE, NONE
ACCESSPATH	ACCESSPATH, NONE
SPACE	SPACE, NONE
NONE	NONE

FORCEROLLUP
Determines if aggregation or rollup of statistics will take place when RUNSTATS is executed in spite of some parts being empty. This will enable the optimizer to select the best access path.

The following are the available options for the FORCEROLLUP Keyword.

YES Indicates that forced aggregation or rollup processing is to be done, even though some parts may not contain data.

NO Indicates that aggregation or rollup will be done only if data is available for all parts.

If data is not available for all parts, DSNU623I message will be issued, provided the installation value for STATISTICS ROLLUP on panel DSNTIPO is set to NO.

KEYCARD
Collects all of the distinct values in all of the 1 to *n* key column combinations for the specified indexes. *n* is the number of columns in the index.

FREQVAL
Controls the collection of frequent value statistics. If you specify FREQVAL, it must be followed by two additional keywords:

NUMCOLS
Indicates the number of key columns to concatenate together when collecting frequent values from the specified index. Specifying '3' means to collect frequent values on the concatenation of the first three key columns. The **default** is **1**, which means collect frequent values on the first key column of the index.

COUNT
Indicates the number of frequent values to be collected. Specifying '15' means collect 15 frequent values from the specified key columns. The **default** is **10**.

NOPAD
Specifies that the variable-length columns in the unloaded or discarded records occupy the actual data length without additional padding. The unloaded records may have varying lengths. If you do not specify NOPAD, default REORG processing pads variable-length columns in the unloaded or discarded records to their maximum length; the unloaded or discarded records have equal lengths for each table.

The NOPAD option can only be specified with UNLOAD EXTERNAL or with UNLOAD DISCARD.

While the LOAD utility processes records with variable-length columns that were unloaded or discarded using the NOPAD option, these records can not be processed by applications that only process fields in fixed positions.

In order for the generated LOAD statement to provide a NULLIF condition for fields that are not in a fixed position, an input field definition is generated with a name in the form of DSN_NULL_IND_*nnnnn*, where *nnnnn* is the number of the associated column.

For example, the LOAD statement generated for the EMP sample table would look similar to the LOAD statement shown in Figure 19.

```
LOAD DATA INDDN SYSREC   LOG NO  RESUME YES
 EBCDIC CCSID(00500,00000,00000)
 INTO TABLE "DSN8710 "."EMP               "
 WHEN(00004:00005 = X'0012')
 ( "EMPNO             " POSITION(00007:00012) CHAR(006)
 , "FIRSTNME          " POSITION(00013)       VARCHAR
 , "MIDINIT           " POSITION(*)           CHAR(001)
 , "LASTNAME          " POSITION(*)           VARCHAR
 ,  DSN_NULL_IND_00005  POSITION(*)           CHAR(1)
 , "WORKDEPT          " POSITION(*)           CHAR(003)
                          NULLIF(DSN_NULL_IND_00005)=X'FF'
 ,  DSN_NULL_IND_00006  POSITION(*)           CHAR(1)
 , "PHONENO           " POSITION(*)           CHAR(004)
                          NULLIF(DSN_NULL_IND_00006)=X'FF'
 ,  DSN_NULL_IND_00007  POSITION(*)           CHAR(1)
 , "HIREDATE          " POSITION(*)           DATE EXTERNAL
                          NULLIF(DSN_NULL_IND_00007)=X'FF'
 ,  DSN_NULL_IND_00008  POSITION(*)           CHAR(1)
 , "JOB               " POSITION(*)           CHAR(008)
                          NULLIF(DSN_NULL_IND_00008)=X'FF'
 ,  DSN_NULL_IND_00009  POSITION(*)           CHAR(1)
```

Figure 19. Sample LOAD statement generated by REORG TABLESPACE with the NOPAD keyword (Part 1 of 2)

```
 , "EDLEVEL           " POSITION(*)           SMALLINT
                          NULLIF(DSN_NULL_IND_00009)=X'FF'
 ,  DSN_NULL_IND_00010  POSITION(*)           CHAR(1)
 , "SEX               " POSITION(*)           CHAR(001)
                          NULLIF(DSN_NULL_IND_00010)=X'FF'
 ,  DSN_NULL_IND_00011  POSITION(*)           CHAR(1)
 , "BIRTHDATE         " POSITION(*)           DATE EXTERNAL
                          NULLIF(DSN_NULL_IND_00011)=X'FF'
 ,  DSN_NULL_IND_00012  POSITION(*)           CHAR(1)
 , "SALARY            " POSITION(*)           DECIMAL
                          NULLIF(DSN_NULL_IND_00012)=X'FF'
 ,  DSN_NULL_IND_00013  POSITION(*)           CHAR(1)
 , "BONUS             " POSITION(*)           DECIMAL
                          NULLIF(DSN_NULL_IND_00013)=X'FF'
 ,  DSN_NULL_IND_00014  POSITION(*)           CHAR(1)
 , "COMM              " POSITION(*)           DECIMAL
                          NULLIF(DSN_NULL_IND_00014)=X'FF'
 )
```

Figure 19. Sample LOAD statement generated by REORG TABLESPACE with the NOPAD keyword (Part 2 of 2)

PUNCHDDN *ddname*

Specifies the DD statement for a data set to receive the LOAD utility control statements that are generated by REORG TABLESPACE UNLOAD EXTERNAL or REORG TABLESPACE DISCARD FROM TABLE ... WHEN.

ddname is the DD name.

The **default** is **SYSPUNCH**.

PUNCHDDN is required if the last partition of a partitioned table space has had its limit key reduced.

The **PUNCHDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

DISCARDDN *ddname*
Specifies the DD statement for a discard data set, which is to hold copies of records that meet the DISCARD FROM TABLE ... WHEN specification.

ddname is the DD name.

If you omit the DISCARDDN option, the utility application program saves discarded records only if a SYSDISC DD statement is in the JCL input.

The **default** is **SYSDISC**.

The **DISCARDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

DISCARD
Specifies that records meeting the specified WHEN conditions are to be discarded during REORG TABLESPACE UNLOAD CONTINUE or UNLOAD PAUSE. If you specify DISCARDDN or a SYSDISC DD statement in the JCL input, discarded records are saved in the associated data set.

DISCARD is valid only for SHRLEVEL NONE or SHRLEVEL REFERENCE. Discarding rows from a table that is part of a referential integrity set sets the CHECK pending status.

Do not specify DISCARD with the UNLOAD EXTERNAL or UNLOAD ONLY option.

FROM TABLE
The table space that is specified in REORG TABLESPACE can store more than one table. All tables are unloaded for UNLOAD EXTERNAL, and all tables might be subject to DISCARD. If you want to qualify the rows to be unloaded or discarded, you must use the FROM TABLE statement.

Do not specify FROM TABLE with REORG TABLESPACE LIST.

table-name
Specifies the name of the table that is to be qualified by the following WHEN clause. The table must be described in the catalog and must not be a catalog table. If the table name is not qualified by an authorization ID, the authorization ID of the person who invokes the utility job step is used as the qualifier of the table name.

WHEN
The WHEN clause tells which records in the table space are to be unloaded (for UNLOAD EXTERNAL) or discarded (for DISCARD). If you do not specify a WHEN clause for a table in the table space, all of the records are unloaded (for UNLOAD EXTERNAL), or none of the records is discarded (for DISCARD).

The option following WHEN describes the conditions for UNLOAD or DISCARD of records from a table.

selection condition
A *selection condition* specifies a condition that is true, false, or unknown about a given row. When the condition is true, the row qualifies for UNLOAD or DISCARD. When the condition is false or unknown, the row does not qualify.

The result of a selection condition is derived by application of the specified *logical operators* (AND, OR, NOT) to the result of each specified predicate. If logical operators are not specified, the result of the selection condition is the result of the specified predicate.

Selection conditions within parentheses are evaluated first. If the order of evaluation is not specified by parentheses, AND is applied before OR.

predicate
A *predicate* specifies a condition that is true, false, or unknown about a row.

labeled-duration-expression
A *labeled-duration-expression* specifies an expression that begins with special register CURRENT DATE or special register CURRENT TIMESTAMP (the forms CURRENT_DATE and CURRENT_TIMESTAMP are also acceptable) and optionally contains arithmetic operations of addition or subtraction expressed by a number followed by one of the seven duration keywords: YEARS, MONTHS, DAYS, HOURS, MINUTES, SECONDS, or MICROSECONDS. (The singular form of these keywords is also acceptable: YEAR, MONTH, DAY, HOUR, MINUTE, SECOND, and MICROSECOND.)

Utilities always evaluate a *labeled-duration-expression* as a timestamp and implicitly performs a conversion to a date if the comparison is with a date column.

Incrementing and decrementing CURRENT DATE: The result of adding a duration to a date, or of subtracting a duration from a date, is itself a date. (For the purposes of this operation, a month denotes the equivalent of a calendar page. Adding months to a date, then, is like turning the pages of a calendar, starting with the page on which the date appears.) The result must fall between the dates January 1, 0001 and December 31, 9999 inclusive. If a duration of years is added or subtracted, only the year portion of the date is affected. The month is unchanged, as is the day unless the result would be February 29 of a non-leap-year. Here the day portion of the result is set to 28.

Similarly, if a duration of months is added or subtracted, only months and, if necessary, years are affected. The day portion of the date is unchanged unless the result would be invalid (September 31, for example). In this case the day is set to the last day of the month.

Adding or subtracting a duration of days will affect the day portion of the date, and potentially the month and year.

Date durations, whether positive or negative, can also be added to and subtracted from dates. As with labeled durations, the result is a valid date.

When a positive date duration is added to a date, or a negative date duration is subtracted from a date, the date is incremented by the specified number of years, months, and days.

When a positive date duration is subtracted from a date, or a negative date duration is added to a date, the date is decremented by the specified number of days, months, and years.

Adding a month to a date gives the same day one month later unless that day does not exist in the later month. In that case, the day in the result is set to the last day of the later month. For example, January 28 plus one month gives February 28; one month added to January 29, 30, or 31 results in either February 28 or, for a leap year, February 29. If one or more months is added to a given date and then the same number of months is subtracted from the result, the final date is not necessarily the same as the original date.

The order in which labeled date durations are added to and subtracted from dates can affect the results. When you add labeled date durations to a date, specify them in the order of YEARS + MONTHS + DAYS. When you subtract labeled date durations from a date, specify them in the order of DAYS - MONTHS - YEARS. For example, to add one year and one day to a date, specify:

```
CURRENT DATE + 1 YEAR + 1 DAY
```

To subtract one year, one month, and one day from a date, specify:

```
CURRENT DATE - 1 DAY - 1 MONTH - 1 YEAR
```

Incrementing and decrementing timestamps: The result of adding a duration to a timestamp, or of subtracting a duration from a timestamp, is itself a timestamp. Date and time arithmetic is performed as previously defined, except that an overflow or underflow of hours is carried into the date part of the result, which must be within the range of valid dates.

basic predicate

A *basic predicate* compares a column with a constant. If the value of the column is null, the result of the predicate is unknown. Otherwise, the result of the predicate is true or false.

Predicate	Is true if and only if
column = constant	column is equal to constant or labeled duration expression.
column < > constant	column is not equal to constant or labeled duration expression.
column > constant	column is greater than constant or labeled duration expression.
column < constant	column is less than constant or labeled duration expression.
column > = constant	column is greater than or equal to constant or labeled duration expression.
column < = constant	column is less than or equal to constant or labeled duration expression.

Note: The following forms of the comparison operators are also supported in basic and quantified predicates: !=, !<, and !>. In addition, in code pages 437, 819, and 850, the forms ¬=, ¬<, and ¬> are supported. All these product-specific forms of the comparison operators are intended only to support existing SQL statements that use these operators and are not recommended for use when writing new SQL statements.

A not sign (¬), or the character that must be used in its place in certain countries, can cause parsing errors in statements passed from one DBMS to another. The problem occurs if the statement undergoes character conversion with certain combinations of source and target CCSIDs. To avoid this problem, substitute an equivalent operator for any operator that includes a not sign. For example, substitute '< >' for '¬=', '<=' for '¬>', and '>=' for '¬<'.

For ASCII table spaces, the constant must be specified in hexadecimal.

BETWEEN predicate

The BETWEEN predicate determines whether a given value lies between two other given values specified in ascending order. Each of the predicate's two other forms has an equivalent search condition, as shown below:

The predicate:	*column* BETWEEN *value1* AND *value2*
is equivalent to:	(*column* >= *value1* AND *column* <= *value2*)
The predicate:	*column* NOT BETWEEN *value1* AND *value2*
is equivalent to:	NOT(*column* BETWEEN *value1* AND *value2*)
and therefore also to:	(*column* < *value1* OR *column* > *value2*)

The values can be constants or labeled duration expressions.

For example, the following predicate is true for any row when salary is greater than or equal 10000 and less than or equal to 20000:

```
SALARY BETWEEN 10000 AND 20000
```

IN predicate

The IN predicate compares a value with a set of values. In the IN predicate, the second operand is a set of one or more values specified by constants.

The predicate:	*value1* IN (*value1*, *value2*,..., *valuen*)
is equivalent to:	(*value1* = *value2* OR ... OR *value1* = *valuen*)
The predicate:	*value1* NOT IN (*value1*, *value2*,..., *valuen*)
is equivalent to:	(*value1* ¬= *value2* AND ... AND *value1* ¬= *valuen*)

For example, the following predicate is true for any row whose employee is in department D11, B01, or C01:

```
WORKDEPT IN ('D11', 'B01', 'C01')
```

LIKE predicate

The *LIKE predicate* qualifies strings that have a certain pattern. Specify the pattern with a string in which the underscore and percent sign characters have special meanings.

Let *x* denote the column to be tested and *y* denote the pattern in the string constant.

The following rules apply to predicates of the form "*x* LIKE *y*...". If NOT is specified, the result is reversed.

- When *x* and *y* are both neither empty nor null, the result of the predicate is true if *x* matches the pattern in *y* and false if *x* does not match the pattern in *y*. Matching the pattern is described below.
- When *x* or *y* is null, the result of the predicate is unknown.
- When *y* is empty and *x* is not, the result of the predicate is false.
- When *x* is empty and *y* is not, the result of the predicate is false unless *y* consists only of one or more percent signs.
- When *x* and *y* are both empty, the result of the predicate is true.

The pattern string and the string to be tested must be of the same type, that is, both *x* and *y* must be character strings or both *x* and *y* must be graphic strings. When *x* and *y* are graphic strings, a character is a DBCS character. When *x* and *y* are character strings and *x* is not mixed data, a character is a SBCS character and *y* is interpreted as SBCS data regardless of is subtype. The rules for **mixed data patterns** are described on page 354.

Within the pattern, a percent sign or underscore character can have a special meaning, or it can represent the literal occurrence of a percent sign or underscore character. To have its literal meaning, it must be preceded by an *escape character*. If it is not preceded by an escape character, it has its special meaning.

The ESCAPE clause designates a single character. That character— and only that character— can be used multiple times within the pattern as an escape character. When the ESCAPE clause is omitted, no character serves as an escape character, so that percent signs and underscores in the pattern always have their special meanings.

The following rules apply to the use of the ESCAPE clause:

- The ESCAPE clause cannot be used if *x* is mixed data.
- If *x* is a character string, the data type of the string constant must be character string. If *x* is a graphic string, the data type of the string constant must be graphic string. In both cases, the length of the string constant must be 1.
- The pattern must not contain the escape character except when followed by the escape character, '%' or '_'. For example, if '+' is the escape character, any occurrences of '+' other than '++', '+_', or '+%' in the pattern is an error.

When that pattern does not include escape characters, a simple description of its meaning is:

- The underscore character (_) represents a single arbitrary character.

- The percent sign (%) represents a string of zero or more arbitrary characters.
- Any other character represents a single occurrence of itself.

More rigorous description of strings and patterns

The string *y* is interpreted as a sequence of the minimum number of substring specifiers such that each character of *y* is part of exactly one substring specifier. A substring specifier is an underscore, a percent sign, or any non-empty sequence of characters other than an underscore or percent sign.

The string *x* matches the pattern *y* if a partitioning of *x* into substrings exists, such that:

- A substring of *x* is a sequence of zero or more contiguous characters and each character of *x* is part of exactly one substring.
- If the *n*th substring specifier is an underscore, the *n*th substring of *x* is any single character.
- If the *n*th substring specifier is a percent sign, the *n*th substring of *x* is any sequence of zero or more characters.
- If the *n*th substring specifier is neither an underscore nor a percent sign, the *n*th substring of *x* is equal to that substring specifier and has the same length as that substring specifier.
- The number of substrings of *x* is the same as the number of substring specifiers.

When escape characters are present in the pattern string, an underscore, percent sign, or escape character represents a single occurrence of itself if and only if it is preceded by an odd number of successive escape characters.

Mixed data patterns: If *x* is mixed data, the pattern is assumed to be mixed data, and its special characters are interpreted as follows:

- A single-byte underscore refers to one single-byte character; a double-byte underscore refers to one double-byte character.
- A percent sign, either single-byte or double-byte, refers to any number of characters of any type, either single-byte or double-byte.
- Redundant shift bytes in *x* or *y* are ignored.

NULL predicate

The *NULL predicate* tests for null values.

If the value of the column is null, the result is true. If the value is not null, the result is false. If NOT is specified, the result is reversed.

reorg tablespace options

For the descriptions of keywords and parameters included within *reorg tablespace options*, see page 355.

REORG TABLESPACE options syntax

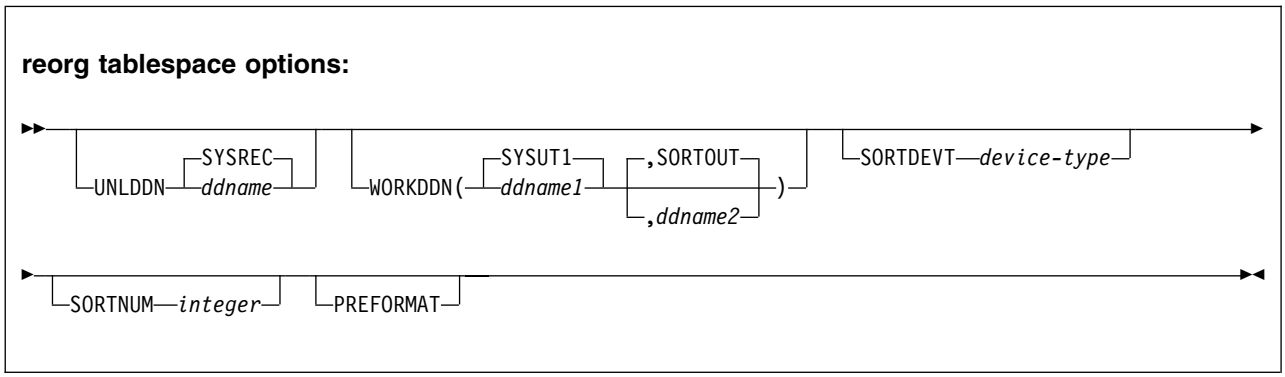

Option descriptions for REORG TABLESPACE options

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

UNLDDN *ddname*

Specifies the DD name of the unload data set.

ddname is the DD name of the unload data set.

The **default** is **SYSREC**.

The **UNLDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

WORKDDN(*ddname1***,***ddname2***)**

ddname specifies the DD statement for a temporary data set used for intermediate output.

ddname1

Is the DD name of the temporary work file for sort input. DB2 requires a work data set for sort input for tables with indexes, unless you specify SORTKEYS.

The **default** is **SYSUT1**.

ddname2

Is the DD name of the temporary work file for sort output. DB2 requires a work data set for sort output, unless you specify SORTKEYS.

The **default** is **SORTOUT**.

Even though WORKDDN is an optional keyword, a DD card for the sort output data set is required in the JCL unless you specify SORTKEYS. If you do not specify WORKDDN, or if you specify it without a *ddname2*, the JCL must have a DD card with the name SORTOUT. If *ddname2* is given, then a DD card must be supplied that matches it.

WORKDDN is ignored for the catalog and directory table spaces listed in "Reorganizing the catalog and directory" on page 369.

The **WORKDDN** keyword specifies either a DDNAME or a TEMPLATE name specification from a previous TEMPLATE control statement. If utility processing detects that the specified name is both a DDNAME in the current job step and a TEMPLATE name, the utility uses DDNAME. For more information about TEMPLATE specifications, see "Chapter 28. TEMPLATE" on page 463.

SORTDEVT *device-type*
Specifies the device type for temporary data sets to be dynamically allocated by DFSORT.

device-type is the device type, and can be any acceptable to the DYNALLOC parameter of the SORT or OPTION control statement for DFSORT.

If you omit SORTDEVT and require a sort of the index keys, you must provide the DD statements that the sort program needs for the temporary data sets.

SORTDEVT is ignored for the catalog and directory table spaces listed in "Reorganizing the catalog and directory" on page 369.

The utility does not allow a TEMPLATE specification to dynamically allocate sort work data sets. The SORTDEVT keyword controls dynamic allocation of these data sets.

SORTNUM *integer*
Specifies the number of temporary data sets to be dynamically allocated by the sort program.

integer is the number of temporary data sets.

If you omit SORTDEVT, SORTNUM is ignored. If you use SORTDEVT and omit SORTNUM, no value is passed to DFSORT. It is allowed to take its own default.

SORTNUM is ignored for the catalog and directory table spaces listed in "Reorganizing the catalog and directory" on page 369.

PREFORMAT
Specifies that the remaining pages are preformatted up to the high allocated RBA in the table space and index spaces associated with the table specified in *table-name*. The preformatting occurs after the data has been loaded and the indexes are built.

PREFORMAT can operate on an entire table space and its index spaces, or on a partition of a partitioned table space and its corresponding partitioning index space.

PREFORMAT is ignored if you specify UNLOAD ONLY or UNLOAD EXTERNAL.

For more information on PREFORMAT, see "Improving performance with LOAD or REORG PREFORMAT" on page 197.

Instructions for running REORG TABLESPACE

To run REORG TABLESPACE, you must:

1. Read "Before running REORG TABLESPACE" on page 357 in this chapter.
2. Prepare the necessary data sets, as described in "Data sets used by REORG TABLESPACE" on page 361.

3. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for REORG TABLESPACE, see "Sample control statements" on page 388.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 365. (For a complete description of the syntax and options for REORG TABLESPACE, see "Syntax and options of the control statement" on page 329.)
5. Check the compatibility table in "Concurrency and compatibility" on page 383 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the REORG TABLESPACE job doesn't complete, as described in "Terminating or restarting REORG TABLESPACE" on page 380.
7. Run REORG TABLESPACE.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Before running REORG TABLESPACE

Catalog and directory table spaces: Before running REORG on a catalog or directory table space, you must take an image copy. Be aware that for the DSNDB06.SYSCOPY, DSNDB01.DBD01, and DSNDB01. SYSUTILX catalog table spaces, REORG scans logs to verify that an image copy is available. If the scan of logs does not find an image copy, DB2 will request archive logs.

Region size: The recommended minimum region size is 4096K.

Mapping table and SHRLEVEL CHANGE: Before running REORG TABLESPACE with SHRLEVEL CHANGE, you must create a **mapping table** and an index for it. The table space that contains the mapping table must be segmented and cannot be the table space to be reorganized. To create the mapping table, use a CREATE TABLESPACE statement similar to the following:

```
CREATE TABLESPACE table-space-name SEGSIZE integer
```

The number of rows in the mapping table should not exceed 110% of the number of rows in the table space or partition to be reorganized. The mapping table must have only the columns and the index created by the following SQL statements:

```
CREATE TABLE table-name1
 (TYPE        CHAR(1) NOT NULL,
  SOURCE_RID  CHAR(5) NOT NULL,
  TARGET_XRID CHAR(9) NOT NULL,
  LRSN        CHAR(6) NOT NULL);
CREATE UNIQUE INDEX index-name1 ON table-name1
 (SOURCE_RID ASC, TYPE, TARGET_XRID, LRSN);
```

The TARGET_XRID column must be specified as CHAR(9) even though the RIDs are still 5 bytes long.

You must have DELETE, INSERT, SELECT, and UPDATE authorization on the mapping table.

You can run more than one REORG SHRLEVEL CHANGE job concurrently, either on separate table spaces or different partitions of the same table space. When running concurrently, each REORG must have a separate mapping table. The mapping tables need not reside in separate table spaces. If only one mapping table

exists, the REORG jobs must be scheduled to run serially. If more than one REORG tries to access the same mapping table at the same time, one of the REORGs will fail.

One way to keep the mapping table of the Online REORG separate is to name the mapping table the same as the utility ID. Since utility IDs must be unique, this is one way to make sure the mapping tables are not used by two reorgs running concurrently. This is not a requirement but a way to keep them apart.

For a sample of REORG with SHRLEVEL CHANGE and a sample mapping table and index, see job sample DSNTEJ1 as described in *DB2 Installation Guide*.

User-managed data sets and SHRLEVEL REFERENCE and CHANGE: When you execute the REORG utility with SHRLEVEL REFERENCE or SHRLEVEL CHANGE, the shadow data sets already exist. The name of the shadow data set is catname.DSNDBx.dbname.psname.y0001.Annn, where y is I or J.

To determine what the existing value is, execute the appropriate query against SYSTABLEPART or SYSINDEXPART:

```
SELECT DBNAME, TSNAME, IPREFIX
  FROM SYSIBM.SYSTABLEPART
  WHERE DBNAME = 'dbname' AND TSNAME = 'psname';
SELECT DBNAME, IXNAME, IPREFIX
  FROM SYSIBM.SYSINDEXES X, SYSIBM.SYSINDEXPART Y
  WHERE X.NAME = Y.IXNAME AND X.CREATOR = Y.IXCREATOR
  AND X.DBNAME = 'dbname' AND X.INDEXSPACE = 'psname';
```

If it is a partitioned table space, you get several rows back from which you select the row for the partitions that you want to reorganize.

The names have the form `catname.DSNDBx.dbname.psname.x0001.Annn`, where x is I or J. Define the data sets as LINEAR and use SHAREOPTIONS(3,3). An efficient method for defining shadow data sets specifies MODEL, so the shadow is created like the original. For example:

```
DEFINE CLUSTER +
       (NAME('catname.DSNDBC.dbname.psname.x0001.A001') +
       MODEL('catname.DSNDBC.dbname.psname.y0001.A001')) +
      DATA                                               +
       (NAME('catname.DSNDBD.dbname.psname.x0001.A001') +
       MODEL('catname.DSNDBD.dbname.psname.y0001.A001') )
```

For user-managed databases, the user must pre-allocate the shadow objects with the correct instance qualifier value prior to executing REORG SHRLEVEL CHANGE or REFERENCE. Assume that the user has a ten-partition table space and wants to determine a naming convention for the data set to successfully executing a REORG SHRLEVEL CHANGE PART 2:6. The following query of the DB2 catalog tables, SYSTABLEPART and SYSINDEXPART, would provide the required information.

```
SELECT DBNAME, TSNAME, PARTITION, IPREFIX FROM SYSIBM.SYSTABLEPART
  WHERE DBNAME = 'DBDV0701' AND TSNAME = 'TPDV0701'
  ORDER BY PARTITION;
SELECT IXNAME, PARTITION, IPREFIX FROM SYSIBM.SYSINDEXPART
  WHERE IXNAME = 'IXDV0701
  ORDER BY PARTITION;
```

The above query would produce the following information.

DBNAME	TSNAME	PARTITION	IPREFIX
DBDV0701	TPDV0701	1	I

DBNAME	TSNAME	PARTITION	IPREFIX
DBDV0701	TPDV0701	4	I
DBDV0701	TPDV0701	3	J
DBDV0701	TPDV0701	2	I
DBDV0701	TPDV0701	5	J
DBDV0701	TPDV0701	6	J
DBDV0701	TPDV0701	7	I
DBDV0701	TPDV0701	8	I
DBDV0701	TPDV0701	9	I
DBDV0701	TPDV0701	10	I

IXNAME	PARTITION	IPREFIX
IXDV0701	10	I
IXDV0701	9	I
IXDV0701	8	I
IXDV0701	7	I
IXDV0701	6	J
IXDBV0701	5	J
IXDV0701	4	I
IXDV0701	3	J
IXDV0701	2	I
IXDV0701	1	I

To execute a REORG SHRLEVEL CHANGE PART 2:6, the following shadow objects must be pre-allocated.

```
vcatnam.DSNDBC.DBDV0701.TPDV0701.J0001.A002
vcatnam.DSNDBC.DBDV0701.TPDV0701.I0001.A003
vcatnam.DSNDBC.DBDV0701.TPDV0701.J0001.A004
vcatnam.DSNDBC.DBDV0701.TPDV0701.I0001.A005
vcatnam.DSNDBC.DBDV0701.TPDV0701.I0001.A006
vcatnam.DSNDBC.DBDV0701.IXDV0701.J0001.A002
vcatnam.DSNDBC.DBDV0701.IXDV0701.I0001.A003
vcatnam.DSNDBC.DBDV0701.IXDV0701.J0001.A004
vcatnam.DSNDBC.DBDV0701.IXDV0701.I0001.A005
vcatnam.DSNDBC.DBDV0701.IXDV0701.I0001.A006
```

If a table space, partition, or index resides in DB2-managed data sets and shadow data sets do not already exist when you execute REORG, DB2 creates the shadow data sets. At the end of REORG processing, the DB2-managed shadow data sets are deleted. You can create the shadows ahead of time for DB2-managed data sets, and it is strongly recommended that you do so for the shadow data set of the logical partition of nonpartitioning indexes.

Estimating the size of pre-allocated data sets: If you have not changed the value of FREEPAGE or PCTFREE, the amount of space required for a shadow data set should be comparable to the amount of space required for the original data set. However, for REORG PART, the space required for the shadow of the logical partition of a nonpartitioning index is approximately equal to the percentage of space the partition occupies in the entire table space. For example, a partitioned

table space with 100 partitions and data relatively evenly balanced across the partitions needs a shadow for the logical partition roughly 1 percent the size of the original nonpartitioning index.

Pre-allocating shadow data sets for REORG PART: By creating the shadow data sets for REORG PART ahead of time, even for DB2-managed data sets, you prevent possible over-allocation of DASD during REORG processing. When reorganizing a partition, you must create the shadow data sets for the partition of the table space and the partition of the partitioning index. In addition, before executing REORG PART with SHRLEVEL REFERENCE or SHRLEVEL CHANGE on partition *mmm* of a partitioned table space, you must create, for each nonpartitioning index that resides in user-defined data sets, a shadow data set for a copy of the logical partition of the index. The name for this shadow data set has the form `catname.DSNDBx.dbname.psname.S0mmm.Annn`.

When reorganizing a range of partitions, you must allocate a single shadow data set for each logical partition. Each logical partition within the range specified will be contained in the single shadow data set. The name for this shadow data set must have the form `catname.DSNDBx.dbname.psname.S0mmm.Annn`, where *mmm* is the first partition in the range specification.

Restart pending status and SHRLEVEL CHANGE: If you specify SHRLEVEL CHANGE, REORG drains the write claim class near the end of REORG processing. In a data sharing environment, if a data sharing member fails and that member has restart pending status for a target page set, the drain can fail. You must postpone running REORG with SHRLEVEL CHANGE until all restart pending statuses have been removed. You can use the DISPLAY GROUP command to determine whether a member's status is FAILED. You can use the DISPLAY DATABASE command with the LOCKS option to determine if locks held.

Data sharing considerations for REORG: You must not execute REORG on an object if another DB2 holds retained locks on the object or has long-running noncommitting applications that use the object. You can use the -DISPLAY GROUP command to determine whether a member's status is ″FAILED.″ You can use the -DISPLAY DATABASE command with the LOCKS option to determine if locks are held.

RECOVER-pending and REBUILD-pending status: You cannot reorganize a table space if:

- Any partition or range of partitions of the partitioned table space is in the RECOVER-pending status
- The clustered index is in the REBUILD-pending status, and the data is unloaded by the cluster index method.

Similarly, you cannot reorganize a single table space partition if:

- The partition is in the RECOVER-pending status
- The corresponding partitioning index is in the REBUILD-pending or RECOVER-pending status, and the data is unloaded by the cluster index method.
- The table space is a subset of a range of partitions that are in REORG-pending status; you must reorganize the entire range to reset the restrictive status.

There is one RECOVER-pending restrictive state:

RECP The table space, index space, or partition of a table space or index space is in a RECOVER-pending status. A single logical partition in RECP does not

restrict access to other logical partitions not in RECP. RECP can be reset by recovering only the single logical partition.

There are three REBUILD-pending restrictive states:

RBDP REBUILD-pending status (RBDP) is set on a physical or logical index partition. The individual physical or logical partition is inaccessible and must be rebuilt using the REBUILD INDEX utility.

PSRBD
Page set REBUILD-pending (PSRBD) is set for nonpartitioning indexes. The entire index space is inaccessible and must be rebuilt using the REBUILD utility.

RBDP*
A REBUILD-pending status that is set only on logical partitions of nonpartitioning indexes. The entire index is inaccessible, but is made available again when the affected partitions are rebuilt using the REBUILD INDEX utility.

For information about resetting the REBUILD-pending and RECOVER-pending states, see Table 134 on page 681 and Table 133 on page 680.

CHECK-pending status: If a table space is in both REORG-pending and CHECK-pending status (or auxiliary CHECK-pending status), run REORG first and then CHECK DATA to clear the respective states. Otherwise, if a table space is not in REORG-pending status, you cannot reorganize a table space or range of partitions if the table space or any partition in the range until the CHECK-pending status is cleared. See “CHECK-pending status” on page 678 for more information about resetting the CHECK-pending status.

REORG-pending status: You must allocate a discard data set (SYSDISC) or specify the DISCARDDN option if the last partition of the table space is in REORG-pending status.

Data sets used by REORG TABLESPACE

Table 67 describes the data sets used by REORG TABLESPACE. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 67. Data sets used by REORG TABLESPACE

Data Set	Description	Required?
SYSDISC	Contains discarded records from REORG DISCARD; optional for REORG The default DD name is SYSDISC.	DISCARDDN
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
SYSPUNCH	Contains a LOAD statement that is generated by REORG, which loads records that REORG DISCARD or REORG UNLOAD EXTERNAL wrote to the DISCARD or UNLOAD data sets. The default DD name is SYSPUNCH.	PUNCHDDN

Table 67. Data sets used by REORG TABLESPACE (continued)

Data Set	Description	Required?
Unload data set	Data set for the unloaded data and the data set to be loaded by the RELOAD phase. The data set is identified by the DD statement or template named in the UNLDDN keyword or by the RECDSN field on the DB2I Utilities Panel. The data set must be a sequential data set that is readable by BSAM. The default DD name is SYSREC. The unload data set must be large enough to contain all the unloaded records from all the tables in the target table space. If there is at least one table in the table space that does not have a clustering index, REORG cannot use the SORTDATA method with SHRLEVEL CHANGE. As a result, the data has to be unloaded in the SYSREC data set.	Yes[2]
Copies	From 1 to 4 output data sets to contain the image copies. Their DD or template names are specified with the COPYDDN and RECOVERYDDN options of the utility control statement.	Yes[3]
Work data sets	Temporary data sets for sort input and output. The DD names have the form DATAWK*nn*.	Yes[4]
Work data sets	Temporary data sets for sort input and output. The DD names have the form SORTWK*nn*.	Yes[5]
Work data sets	Two temporary data sets for sort input and sort output. Their DD or template names are specified with the WORKDDN option of the utility control statement. The default DD name for sort input is SYSUT1. The default DD name for sort output is SORTOUT.	No[1]

Note:

[1] Not required if SORTKEYS is used; otherwise, required for tables with indexes.

[2] Required unless NOSYSREC or SHRLEVEL CHANGE is specified.

[3] Required if COPYDDN, RECOVERYDDN, SHRLEVEL REFERENCE, or SHRLEVEL CHANGE is specified.

[4] Required if NOSYSREC or SHRLEVEL CHANGE is specified but SORTDEVT is not specified.

[5] Required if any indexes exist and SORTDEVT is not specified.

The following objects are named in the utility control statement and do not require DD cards in the JCL:

Table space
The name of the table space to be reorganized. It is named in the control statement and is accessed through the DB2 catalog.

Calculating the size of the unload data set: The size required for the unload data set varies depending on the options used for REORG.

1. If you use REORG with UNLOAD PAUSE or CONTINUE and with KEEPDICTIONARY (assuming a compression dictionary already exists), the size of the unload data set, in bytes, can be roughly calculated as the VSAM hi-used RBA for the table space. The hi-used RBA can be obtained from the associated VSAM catalog. For SHRLEVEL CHANGE, add (number of records * 11) bytes to the VSAM hi-used RBA.
2. If you use REORG UNLOAD ONLY, or UNLOAD PAUSE or CONTINUE without KEEPDICTIONARY, the size of the unload data set, in bytes, can be calculated as the maximum row length multiplied by the number of rows. The maximum row length is the row length, including the 6 byte record prefix, plus the length of the longest clustering key. If there are multiple tables in the table space, the formula is:

   ```
   Sum over all tables (row length × number of rows)
   ```

   For SHRLEVEL CHANGE, also add:

   ```
   (21 × ((NEARINDREF + FARINDREF) × 1.1))
   ```

   where:

   NEARINDREF Value obtained from the NEARINDREF column of the SYSIBM.SYSTABLEPART catalog table[4].

   FARINDREF Value obtained from the FARINDREF column of the SYSIBM.SYSTABLEPART catalog table[4].
3. If you have variable length fields, the calculation in 2 might give you excessive space. Use the average uncompressed row length multiplied by the number of rows.

For certain table spaces in the catalog and directory, the unload data set for the table spaces will have a different format. The calculation for the size of this data set is as follows:

```
data set size in bytes = (28 + longrow) × numrows
```

where:

longrow Length of the longest row in the table space

numrows The number of rows in the data set

The length of the row is calculated:

```
Sum of column lengths + 4 bytes for each link
```

The length of the column is calculated:

```
Maximum length of the column + 1 (if nullable) + 2 (if varying
  length)
```

See “Reorganizing the catalog and directory” on page 369 for more information about reorganizing catalog and directory table spaces.

4. The accuracy of the data set size calculation depends on recent information in the SYSTABLEPART catalog table.

Calculating the size of the work data sets: When reorganizing an index space or a table space with indexes, you need a non-DB2 sequential work data set unless you specify the SORTKEYS keyword. That data set is identified by the DD statement named in the WORKDDN option. During the RELOAD phase, the index keys and the data pointers are unloaded to the work data set. This data set is used to update the index data pointers after the data has been moved. It is required only during the execution of REORG.

To calculate the approximate size (in bytes) of both WORKDDN data sets SORTOUT and SYSUT1, follow these steps:

1. For each table, calculate the number of keys:

   ```
   number of keys = (#tablerows × #indexes)
   ```

   where:

   #tablerows Number of records in the table.

   #indexes Number of indexes defined on the table.

   For SHRLEVEL CHANGE, *#indexes* should count the number of indexes on the table, plus 1 for the mapping index.
2. Add the sums obtained in step 1.

   For SHRLEVEL CHANGE, also add:

   ```
   ((NEARINDREF + FARINDREF) × 1.1)
   ```

   where:

   NEARINDREF Value obtained from the NEARINDREF column of the SYSIBM.SYSTABLEPART catalog table[4].

   FARINDREF Value obtained from the FARINDREF column of the SYSIBM.SYSTABLEPART catalog table[4].
3. Multiply the sum in step 2 by the longest key length plus 9. When determining the longest key length, remember that the length of the mapping index is 21 bytes.

Allocating twice the space used by the input data sets is usually adequate for the sort work data sets. For compressed data, double again the space allocated for the sort work data sets if you use the following REORG options:

- UNLOAD PAUSE without KEEPDICTIONARY
- UNLOAD CONTINUE without KEEPDICTIONARY

Two or three large SORTWK*nn* data sets are preferable to several small ones. If adequate space is not available, you cannot run REORG.

Specifying a destination for DFSORT messages: The REORG utility job step must contain a UTPRINT DD statement to define a destination for messages issued by DFSORT during the SORT phase of REORG. The default DD statement used by DB2I and the %DSNU CLIST command and by the DSNUPROC procedure is:

```
//UTPRINT DD  SYSOUT=A
```

Creating the control statement

See “Syntax and options of the control statement” on page 329 for REORG TABLESPACE syntax and option descriptions. See “Sample control statements” on page 388 for examples of REORG TABLESPACE usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

"Determining when an object should be reorganized"
"Specifying access with SHRLEVEL" on page 367
"Omitting the output data set" on page 368
"Unloading without reloading" on page 368
"Reclaiming space from dropped tables" on page 369
"Considerations for fallback recovery" on page 369
"Reorganizing the catalog and directory" on page 369
"Changing data set definitions" on page 371
"Temporarily interrupting REORG" on page 371
"Building a compression dictionary" on page 372
"Overriding dynamic DFSORT and SORTDATA allocation" on page 372
"Rebalancing partitions using REORG" on page 372
"Using inline COPY with REORG TABLESPACE" on page 373
"Improving performance" on page 373
"Improving performance with LOAD or REORG PREFORMAT" on page 197
"Building indexes in parallel for REORG TABLESPACE" on page 375

Determining when an object should be reorganized

Product-sensitive Programming Interface

You can determine when to run REORG for non-LOB table spaces and indexes by using the OFFPOSLIMIT, INDREFLIMIT catalog query options. If you specify the REPORTONLY option, REORG will produce a report detailing if a REORG is recommended; a REORG is not performed.

When you specify the catalog query options along with the REPORTONLY option, REORG produces a report with one of the following return codes:

1 No limit met; no REORG performed or recommended.
2 REORG performed or recommended.

Alternatively, information from the SYSTABLEPART and SYSINDEXPART catalog tables can tell you which table spaces and indexes qualify for reorganization. This information can also be used to determine when the DB2 catalog table spaces require reorganization. For table spaces SYSDBASE, SYSVIEWS, and SYSPLAN of the catalog, the value for columns FAROFFPOS and NEAROFFPOS of SYSINDEXPART should not be used when determining whether to reorganize.

Table spaces or partitions that are in REORG-pending status should be reorganized. Use the DISPLAY DATABASE RESTRICT command to display those table spaces and partitions that require reorganization. See "Appendix C. Resetting an advisory or restrictive status" on page 677 for more information.

Information from the SYSTABLEPART catalog table can also tell you how well DASD space is being used. If you want to find the number of varying-length rows relocated to other pages because of an update, run RUNSTATS and issue this statement:

```
SELECT CARD, NEARINDREF, FARINDREF
  FROM SYSIBM.SYSTABLEPART
  WHERE DBNAME = 'XXX'
  AND TSNAME = 'YYY';
```

A large number (relative to previous values you have received) for FARINDREF indicates that I/O activity on the table space is high. If you find that this number increases over a period of time, you probably need to reorganize the table space to improve performance, and increase PCTFREE or FREEPAGE for the table space with the ALTER TABLESPACE statement.

The following statement returns the percentage of unused space in nonsegmented table space YYY. In nonsegmented table spaces, the space used by dropped tables is not reclaimed until you reorganize the table space.

```
SELECT PERCDROP
  FROM SYSIBM.SYSTABLEPART
  WHERE DBNAME = 'XXX'
  AND TSNAME = 'YYY';
```

Issue the following statement to determine whether the rows of a table are stored in the same order as the entries of its clustering index (XXX.YYY):

```
SELECT NEAROFFPOSF, FAROFFPOSF
  FROM SYSIBM.SYSINDEXPART
  WHERE IXCREATOR = 'index_creator_name'
  AND IXNAME = 'index_name';
```

There are several indicators available to signal a time for reorganizing the table spaces. A large value for FAROFFPOSF might indicate that clustering is degenerating. Reorganizing the table space would improve performance.

A large value for NEAROFFPOSF might indicate also that reorganization might improve performance. However, in general it is not as critical a factor as FAROFFPOSF.

FAROFFPOSF and NEAROFFPOSF do not have performance considerations for certain DB2 catalog tables.

DSNDB06.SYSDBASE
DSNDB06.SYSDBAUT
DSNDB06.SYSGROUP
DSNDB06.SYSPLAN
DSNDB06.SYSVIEWS
DSNDB01.DBD01

For any table, the REORG utility repositions rows into the sequence of the key of the clustering index defined on that table. If you specify the SORTDATA option of the REORG utility, the data is unloaded using a sequential scan. If you do not specify the SORTDATA option, REORG uses the clustering index to unload the data.

For nonclustering indexes, the statistical information recorded by RUNSTATS in SYSINDEXES and SYSINDEXPART might appear even worse after the clustering index is used to reorganize the data. This applies only to CLUSTERING and CLUSTERED in SYSINDEXES and to NEAROFFPOS and FAROFFPOS in SYSINDEXPART.

In general, it is a good practice to run RUNSTATS if the statistics are not current. If you have an object that should also be reorganized, run REORG with STATISTICS and take inline copies. If you run REORG PART and nonpartitioning indexes exist, subsequently run RUNSTATS for each nonpartitioning index.

End of Product-sensitive Programming Interface

Specifying access with SHRLEVEL

For reorganizing a table space, or a partition of a table space, the SHRLEVEL option lets you choose the level of access you have to your data during reorganization:

- REORG with SHRLEVEL NONE, the default, reloads the reorganized data into the original area being reorganized. Applications have read-only access during unloading and no access during reloading. SHRLEVEL NONE is the only access level that resets REORG-pending status.
- REORG with SHRLEVEL REFERENCE reloads the reorganized data into a new (shadow) copy of the area being reorganized. Near the end of reorganization, DB2 switches applications' future access from the original to the shadow copy. For SHRLEVEL REFERENCE, applications have read-only access during unloading and reloading, and a brief period of no access during switching.
- REORG with SHRLEVEL CHANGE reloads the reorganized data into a shadow copy of the area being reorganized. For REORG TABLESPACE SHRLEVEL CHANGE, a mapping table maps between RIDs in the original copy of the table space or partition and RIDs in the shadow copy; see page 357 for instructions on creating the mapping table. Applications can read and write the original area, and DB2 records the writing in the log. DB2 then reads the log and applies it to the shadow copy to bring the shadow copy up to date. This step executes iteratively, with each iteration processing a sequence of log records. Near the end of reorganization, DB2 switches applications' future access from the original to the shadow copy. Applications have read/write access during unloading and reloading, a brief period of read-only access during the last iteration of log processing, and a brief period of no access during switching.

Log processing with SHRLEVEL CHANGE: When you specify SHRLEVEL CHANGE, DB2 processes the log to update the shadow copy. This step executes iteratively. The first iteration processes the log records that accumulated during the previous iteration. The iterations continue until one of these conditions is met:

- DB2 estimates that the time to perform the log processing in the next iteration will be less than or equal to the time specified by MAXRO. If this condition is met, the next iteration will be the last.
- DB2 estimates that the SWITCH phase will not start by the deadline specified by DEADLINE. If this condition is met, DB2 terminates reorganization.
- The number of log records that the next iteration will process is not sufficiently lower than the number of log records processed in the previous iteration. If this condition is met but the first two conditions are not, DB2 sends message DSNU377I to the console. DB2 continues log processing for the length of time specified by DELAY and then performs the action specified by LONGLOG.

Operator actions: LONGLOG specifies the action that DB2 performs if log processing is not catching up. See "Option descriptions" on page 335 for a description of the LONGLOG options. If no action is taken after message DSNU377I is sent to the console, the LONGLOG option automatically goes into effect. Some examples of possible actions you may take:

- Execute the START DATABASE(db) SPACENAM(ts) ... ACCESS(RO) command and the QUIESCE utility to drain the write claim class. DB2 performs the last iteration, if MAXRO is not DEFER. After the QUIESCE, you should also execute the ALTER UTILITY command, even if you do not change any REORG parameters.
- Execute the START DATABASE(db) SPACENAM(ts) ... ACCESS(RO) command and the QUIESCE utility to drain the write claim class. Then, after reorganization has made some progress, execute the START DATABASE(db) SPACENAM(ts) ...

ACCESS(RW) command. This increases the likelihood that log processing will catch up. After the QUIESCE, you should also execute the ALTER UTILITY command, even if you do not change any REORG parameters.

- Execute the ALTER UTILITY command to change the value of MAXRO. Changing it to a huge positive value, such as 9999999, causes the next iteration to be the last iteration.
- Execute the ALTER UTILITY command to change the value of LONGLOG.
- Execute the TERM UTILITY command to terminate reorganization.
- Adjust the amount of buffer space allocated to reorganization and to applications. This can increase the likelihood that log processing will catch up. After adjusting the space, you should also execute the ALTER UTILITY command, even if you do not change any REORG parameters.
- Adjust the scheduling priorities of reorganization and applications. This can increase the likelihood that log processing will catch up. After adjusting the priorities, you should also execute the ALTER UTILITY command, even if you do not change any REORG parameters.

DB2 does not take the action specified in the LONGLOG phrase if any one of these events occurs before the delay expires:

- An ALTER UTILITY command is issued.
- A TERM UTILITY command is issued.
- DB2 estimates that the time to perform the next iteration will be less than or equal to the time specified in the MAXRO phrase.
- REORG terminates for any reason (including the deadline).

Omitting the output data set

For REORG TABLESPACE, you can use the NOSYSREC option to omit the unload data set. You can use this option only if you specify SORTDATA, SHRLEVEL REFERENCE, or SHRLEVEL NONE, and only if you do not specify UNLOAD PAUSE or UNLOAD ONLY. This option provides a performance advantage. However, you should be aware of the following:

- For REORG TABLESPACE SORTDATA NOSYSREC, DB2 assumes there is a clustering index present.
- For REORG TABLESPACE SHRLEVEL CHANGE, REORG omits the unload data set, even if you omit NOSYSREC, unless there is no explicit clustering index.
- For REORG TABLESPACE SHRLEVEL REFERENCE, if you do not use the NOSYSREC option and an error occurs during reloading, you can restart at the RELOAD phase of REORG using the contents of the unload data set. However, if you specify both SORTDATA and NOSYSREC, you must restart at the UNLOAD phase.
- For REORG TABLESPACE SHRLEVEL NONE with NOSYSREC, if an error occurs during reloading, you must execute the RECOVER TABLESPACE utility, starting from the most recent image copy. Therefore, if you specify NOSYSREC with SHRLEVEL NONE, you must create an image copy before starting REORG TABLESPACE in addition to any image copies you create during or after REORG.

Unloading without reloading

REORG can unload data without continuing and without leaving a SYSIBM.SYSUTIL record after the job ends.

If you specify UNLOAD ONLY, REORG unloads data from the table space and then ends. You can reload the data at a later date with the LOAD utility, specifying FORMAT UNLOAD.

Between unloading and reloading, you may add a validation routine to a table. On reloading, all the rows will be checked by the validation procedure.

REORG UNLOAD ONLY should **not** be used for data propagation. When you specify the UNLOAD ONLY option, REORG only unloads what physically resides in the base table space; LOBs are not unloaded. For purposes of data propagation, you should use REORG UNLOAD EXTERNAL instead.

Reclaiming space from dropped tables

Reorganization omits tables that have been previously dropped, reclaiming the space they acquired. See “Reclaiming space in the DBD” on page 235 for actions to take when you drop a table.

Considerations for fallback recovery

If RECOVER cannot use the latest image copy or copies as a starting point for the recovery, it attempts to use previous copies; if that attempt fails, it restores from the log.

However, if you use REORG SHRLEVEL NONE LOG NO, RECOVER cannot restore from the log past the point at which the object was last reorganized successfully. Therefore, you must take an image copy after running REORG with LOG NO to establish a level of fall back recovery.

To create a new recovery point, it is strongly recommended that immediately following an ALTER INDEX operation that modifies key values, you either:

- Run REORG with COPYDDN and SHRLEVEL NONE specified, or
- Take a full image copy immediately after REORG completes.

If you performed a REORG to turn off REORG-pending status (REORP), you should also take an inline image copy, or run the COPY utility. Image copies taken prior to resetting the REORG-pending status cannot be used for recovery to current RBA or LRSN.

Successful REORG LOG NO processing inserts an SYSIBM.SYSCOPY row with ICTYPE=’W’ for each index that was defined with COPY YES. REORG also places a reorganized index in informational COPY-pending status. You should take a full image copy of the index after the REORG job completes to create a valid point of recovery.

Reorganizing the catalog and directory

You can run REORG TABLESPACE on the table spaces in the catalog database
(DSNDB06) and the SCT02, SPT01, DBD01, and SYSLGRNX table spaces in the
directory database (DSNDB01).

Attention: You must take a full image copy before and after reorganizing any catalog or directory object.

When you REORG the DSNDB06.SYSCOPY table space with the LOG NO option and omit the COPYDDN option, DB2 places the table space in COPY-pending status. Take a full image copy of the table space to remove the COPY-pending status before continuing to reorganize the catalog or directory table spaces.

Running REORG LOG NO COPYDDN avoids the COPY-pending status, because an inline copy is taken during the REORG.

The FASTSWITCH YES option is ignored for catalog and directory objects.

When to run REORG on the catalog and directory: You should not need to run REORG TABLESPACE on the catalog and directory table spaces as often as you do on user table spaces. The statistics collected by RUNSTATS that you use to determine if a REORG is required for a user table space can also be used for the catalog table spaces. The only difference is the information in the columns NEAROFFPOS and FAROFFPOS in table SYSINDEXPART. These columns can tolerate a higher value before a reorganization is needed if the table space is DSNDB06.SYSDBASE, DSNDB06.SYSVIEWS, or DSNDB06.SYSPLAN. When it is determined that any of the following catalog table spaces require reorganization, you should also reorganize the corresponding directory table space:

Catalog Table Space	Directory Table Space
DSNDB06.SYSDBASE	DSNDB01.DBD01
DSNDB06.SYSPLAN	DSNDB01.SCT02
DSNDB06.SYSPKAGE	DSNDB01.SPT01

Fragmentation and wasted space in the catalog table spaces affect the performance of user queries against the catalog and performance of DB2 functions.

Associated directory table spaces: When certain catalog table spaces are reorganized, you should reorganize the associated directory table space as well.

Limitations for reorganizing the catalog and directory:

- You cannot reorganize DSNDB01.SYSUTILX or DSNDB01.SYSLGRNGX.
- The UNLOAD ONLY and LOG YES options are not allowed for catalog and directory table spaces. However, LOG YES is required for the catalog LOB table spaces.
- The WORKDDN, SORTDATA, SORTDEVT, SORTNUM, SORTKEYS, COPYDDN, and RECOVERYDDN options are ignored for the following catalog and directory table spaces:
  DSNDB06.SYSDBASE
  DSNDB06.SYSDBAUT
  DSNDB06.SYSGROUP
  DSNDB06.SYSPLAN
  DSNDB06.SYSVIEWS
  DSNDB01.DBD01
- REORG TABLESPACE with SHRLEVEL REFERENCE or CHANGE cannot operate on the following catalog and directory table spaces:
  DSNDB06.SYSDBASE
  DSNDB06.SYSDBAUT

DSNDB06.SYSGROUP
DSNDB06.SYSPLAN
DSNDB06.SYSVIEWS
DSBDB06.SYSSTATS
DSNDB06.SYSHIST
DSNDB01.DBD01

- REORG TABLESPACE with STATISTICS cannot collect inline statistics on the following catalog and directory table spaces:
  DSNDB06.SYSDBASE
  DSNDB06.SYSDBAUT
  DSNDB06.SYSGROUP
  DSNDB06.SYSPLAN
  DSNDB06.SYSVIEWS
  DSBDB06.SYSSTATS
  DSNDB06.SYSHIST
  DSNDB01.DBD01

Phases for reorganizing the catalog and directory: REORG TABLESPACE processes certain catalog and directory table spaces differently from other table spaces; it does not execute the build and sort phases for the following table spaces:

DSNDB06.SYSDBASE
DSNDB06.SYSDBAUT
DSNDB06.SYSGROUP
DSNDB06.SYSPLAN
DSNDB06.SYSVIEWS
DSNDB01.DBD01

For these table spaces, REORG TABLESPACE reloads the indexes (in addition to the table space) during the reload phase, rather than storing the index keys in a work data set for sorting.

Changing data set definitions

If the table space is defined by storage groups, space allocation is handled by DB2 and data set definitions cannot be altered during the reorganization process. DB2 deletes and redefines the necessary data sets to reorganize the object.

For REORG with SHRLEVEL REFERENCE or CHANGE, you can use the ALTER STOGROUP command to change the characteristics of a DB2-managed data set. The user effectively changes the characteristics of a user-managed data set by specifying the desired new characteristics when creating the shadow data set; see page 358 for more information about user-managed data sets. In particular, placing the original and shadow data sets on different DASD volumes might reduce contention and thus improve the performance of REORG and the performance of applications during REORG execution.

Temporarily interrupting REORG

You can temporarily pause REORG. If you specify UNLOAD PAUSE, REORG pauses after unloading the table space into the unload data set. You cannot use NOSYSREC and PAUSE. The job completes with return code 4. You can restart REORG using the phase restart or current restart. The REORG statement must not be altered.

The SYSIBM.SYSUTIL record for the REORG utility remains in ″stopped″ status until REORG is restarted or terminated.

While REORG is interrupted by PAUSE, you may re-define the table space attributes for user defined table spaces. PAUSE is not required for STOGROUP defined table spaces. Attribute changes are done automatically by a REORG following an ALTER TABLESPACE.

Building a compression dictionary

The compression dictionary is built during the UNLOAD phase. This dictionary is then used during the RELOAD phase to compress the data. Specify the KEEPDICTIONARY option to save the cost of rebuilding the dictionary if you are satisfied with the current compression ratio.

Overriding dynamic DFSORT and SORTDATA allocation

When you specify SORTDATA on your REORG statement, DB2 estimates how many rows are to be sorted and passes this information to DFSORT on the parameter FILSZ, letting DFSORT dynamically allocate the necessary sort workspace.

If the table space contains rows with VARCHAR columns, DB2 might not be able to accurately estimate the number of rows. If the estimated number of rows is too high, and the sort work space is not available, DFSORT might fail and cause an abend.

If compression is defined for the table space, REORG doubles the estimated FILSZ so that DFSORT allocates enough space to expand the compressed records during the UNLOAD phase.

You can override this dynamic allocation of sort workspace in two ways:

- Allocate the sort work data sets with SORTWK*nn* DD statements in your JCL
- Override DB2's row estimate in FILSZ using control statements passed to DFSORT. However, using control statements overrides size estimates passed to DFSORT in all invocations of DFSORT in the job step, including sorting keys to build indexes, and any sorts done in any other utility executed in the same step. The result might be reduced sort efficiency or an abend due to an out of space condition.

Rebalancing partitions using REORG

If you use ALTER INDEX to modify the limit keys for partition boundaries, you must subsequently use REORG TABLESPACE to redistribute data in the partitioned table spaces based on the new key values and to reset the REORG-pending status. The following example specifies options that help maximize performance while performing the necessary rebalancing REORG:

```
REORG TABLESPACE DSN8S71E PART 2:3
 SORTDATA
 NOSYSREC
 SORTKEYS
 COPYDDN SYSCOPY
 STATISTICS TABLE INDEX(ALL)
```

You can reorganize a range of partitions, even if the partitions are not in REORG-pending status. If you specify the STATISTICS keyword, REORG collects data on the specified range of partitions.

If you perform a REORG on partitions that are in the REORG-pending status, be aware that:

- You must specify SHRLEVEL NONE if the object is in REORG-pending status. Otherwise, REORG terminates and issues message DSNU273I and return code 8.

- REORG ignores the KEEPDICTIONARY option for any partition that is in REORG pending status; REORG automatically rebuilds the dictionaries for the affected partitions. However, if you specify a range of partitions that includes some partitions that are not in REORG-pending restrictive status, REORG honors the KEEPDICTIONARY option for those non-restricted partitions.
- If any partition is in REORG-pending status when REORG executes, DB2 writes a SYSCOPY record with STYPE='A' for each partition that is specified on the REORG job.
- If you take an inline image copy of a range of partitions, DB2 writes one SYSCOPY record with ICTYPE='F' for each part, and each record has the same data set name.
- Specify the DISCARDDN and PUNCHDDN data sets for a table space that is defined as LARGE or DSSIZE, but has subsequently had an ALTER INDEX statement reduce the limit key for the last partition of the table space. Otherwise, REORG will terminate issuing message DSNU035I and return code 8.

You cannot reorganize a subset of a range of partitions that are in REORG pending status; you must reorganize the entire range to reset the restrictive status.

Using inline COPY with REORG TABLESPACE

You can create a full image copy data set (SHRLEVEL REFERENCE) during REORG TABLESPACE execution. The new copy is an **inline copy**. The advantage to using inline copy is that the table space is not left in COPY-pending status regardless of which LOG option was specified for the utility. Thus, data availability is increased.

To create an inline copy, use the **COPYDDN** and **RECOVERYDDN** keywords. You can specify up to two primary and two secondary copies. Inline copies are produced during the RELOAD phase of REORG processing.

The SYSCOPY record produced by an inline copy contains ICTYPE=F, SHRLEVEL=R. The STYPE column contains an X if the image copy was produced by REORG TABLESPACE LOG(YES), and an W if the image copy was produced by REORG TABLESPACE LOG(NO). The data set produced by the inline copy is logically equivalent to a full image copy with SHRLEVEL REFERENCE, but the data within the data set differs in some respects:

- Data pages might be out of sequence and some might be repeated. If pages are repeated, the last one is always the correct copy.
- Space map pages will be out of sequence and might be repeated

The total number of duplicate pages will be small, with a negligible effect on the space required for the data set. One exception to this guideline is in the case of running REORG SHRLEVEL CHANGE, where the number of duplicate pages will vary with the number of records applied during the LOG phase.

Improving performance

To improve REORG performance:

- We strongly recommend that you specify SORTDATA on your REORG statement unless your data set is very large, and you do not want to allocate the extra DASD required by DFSORT. In this case, DB2 unloads the data by table space scan and executes DFSORT to sort the data into clustering order. SORTDATA is useful if either:
  - Your table's CLUSTERRATIOF is less than 95% or
  - FAROFFPOS/CARD for your table is greater than 5%

In general, the lower the CLUSTERRATIOF of your tables, the greater the performance improvement of REORG when specifying SORTDATA.

- Run REORG concurrently on separate partitions of a partitioned table space. When you run REORG on partitions of a partitioned table space, the sum of each job's processor usage is greater than for a single REORG of the entire table space. However, the elapsed time of reorganizing the entire table in parallel may be significantly less than it would be for a single REORG job.
- Specify SORTKEYS on your REORG statement to sort index keys in parallel with the reload and build phases. This option passes index keys in sort to memory rather than writing them to sort input and output data files. In addition to improving performance, this option reduces work space requirements because the SYSUT1 and SORTOUT data sets are not required. This option is recommended if more than one index needs to be created. REORG will implicitly use SORTKEYS if you specified SHRLEVEL CHANGE. However, if a job using SORTKEYS abends in the RELOAD, SORT, BUILD, or SORTBLD phase, it can only be restarted at the beginning of the reload phase.

  Use parallel index build for table spaces or partitions with more than one index defined. For more information, see "Building indexes in parallel for REORG TABLESPACE" on page 375.
- Specify NOSYSREC on your REORG statement. See "Omitting the output data set" on page 368 for restrictions.
- If you are using 3990 caching, and you have the nonpartitioning indexes on RAMAC, consider specifying YES on the UTILITY CACHE OPTION field of installation panel DSNTIPE. This allows DB2 to use sequential prestaging when reading data from RAMAC for the following utilities:
  - LOAD PART *integer* RESUME
  - REORG TABLESPACE PART

  For these utilities, prefetch reads remain in the cache longer, thus possibly improving performance of subsequent writes.

| Use Inline Copy and Inline Statistics instead of running separate COPY and
| RUNSTATS utilities.

When to use SHRLEVEL CHANGE: Schedule REORG with SHRLEVEL CHANGE when the rate of writing is low and transactions are short. Avoid scheduling REORG with SHRLEVEL CHANGE when low-tolerance applications are executing.

Performance implications with SHRLEVEL CHANGE: Under certain circumstances, the log records used by REORG SHRLEVEL CHANGE contain additional information, as if DATA CAPTURE CHANGES were used. Generation of the additional information can slow applications and increase consumption of log space. The additional information is generated for all the tables in the table space if at least one table satisfies all these conditions:

- The table has undergone ALTER TABLE ADD column
- The table does not use DATA CAPTURE CHANGES
- One of these conditions is true:
  - The area being reorganized uses data compression
  - The area is a partitioned table space, and at least one partition uses data compression

| ***When to use DRAIN_WAIT:*** The DRAIN_WAIT option is provided to give greater
| control over the time online REORG will wait for drains. Also because the

DRAIN_WAIT is the aggregate time that online REORG will wait to perform a drain on a table space and associated indexes, the length of drains will be more predictable than each part and index having it's own individual waiting time limit.

By specifying a short delay time (less than the system timeout value, IRLMRWT), the impact on applications can be reduced by reducing timeouts. The RETRY option could be used to give the online REORG further chances to complete successfully. If the user does not wish to use RETRY processing, DRAIN_WAIT can still be used to set a specific and more consistent limit on the length of drains.

RETRY allows an online REORG that is unable to drain the objects that it requires, to try again after a set period (RETRY_DELAY). During the RETRY_DELAY period, all the objects are available for READ/WRITE in the case of SHRLEVEL CHANGE. For SHRLEVEL REFERENCE, the objects remain with the access prior to the attempted drain (i.e. if the drain fails in the UNLOAD phase the object remains R/W, if the drain fails in the SWITCH phase, then objects remain in READ ONLY). Since application SQL can queue up behind any unsuccessful drain the online REORG has tried, it is recommended to have a reasonable delay before retrying, to allow this work to complete: the default is 5 minutes.

When DRAIN WRITERS (the default) is used with SHRLEVEL CHANGE and RETRY, multiple read only log iterations can occur. Since generally, online REORG may have to do more work when RETRY is specified and this may result in multiple or extended periods of restricted access. It is essential that applications running alongside online REORG, perform frequent commits. During the interval between retries, the utility is still active and consequently other utility activity against the table space and/or indexes are restricted.

When doing a table space REORG with RETRY and SHRLEVEL CHANGE both specified, the size of the COPY taken by REORG may increase.

Recommendation: Run online REORG during lighter periods of activity on the table space or index.

Building indexes in parallel for REORG TABLESPACE: Use parallel index build to reduce the elapsed time for a REORG TABLESPACE job by sorting the index keys and rebuilding multiple indexes in parallel, rather than sequentially. Optimally, a pair of subtasks process each index; one subtask sorts extracted keys while the other subtask builds the index. REORG TABLESPACE begins building each index as soon as the corresponding sort emits its first sorted record. For more information about improving index key sort performance, see "Improving performance" on page 373.

Figure 20 on page 376 shows a REORG TABLESPACE flow with parallel index build, which requires SORTKEYS. DB2 starts multiple subtasks to sort index keys and build indexes in parallel. If you specified STATISTICS, additional subtasks collect the sorted keys and update the catalog table in parallel, eliminating the need for a second scan of the index by a separate RUNSTATS job.

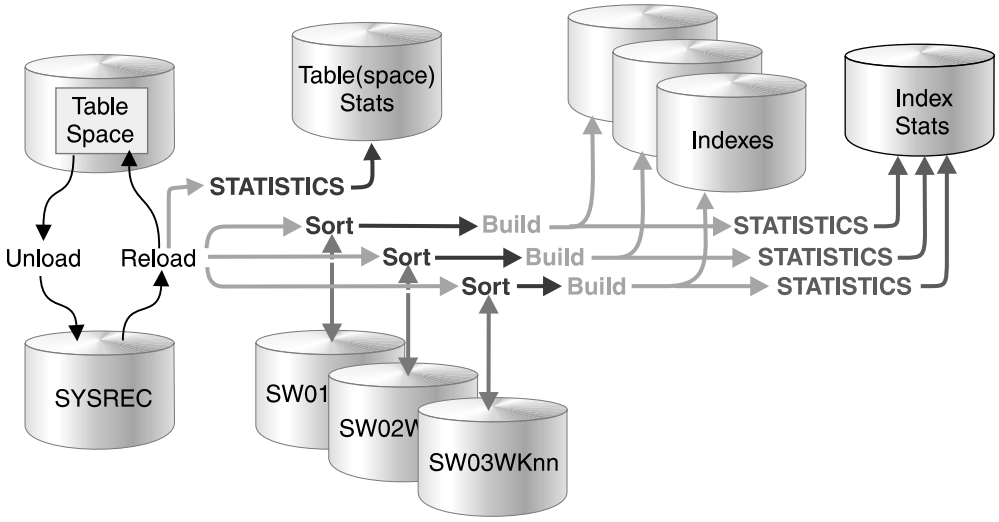

Figure 20. Building indexes using parallel index build

REORG TABLESPACE uses parallel index build if all of the following conditions are true:

- There is more than one index to be built (including the mapping index for SHRLEVEL CHANGE).
- You specify either SORTKEYS or SHRLEVEL CHANGE in the utility statement.
- You either allow the utility to dynamically allocate the data sets needed by SORT, or provide the necessary data sets yourself.

Select one of the following methods to allocate sort work and message data sets:

Method 1: REORG TABLESPACE determines the optimal number of sort work and message data sets.

1. Specify the SORTKEYS and SORTDEVT keywords in the utility statement.
2. Allow dynamic allocation of sort work data sets by *not* supplying SORTWK*nn* DD statements in the REORG TABLESPACE utility JCL.
3. Allocate UTPRINT to SYSOUT.

Method 2: Allows you to control allocation of sort work data sets, while REORG TABLESPACE allocates message data sets.

1. Specify the SORTKEYS keyword in the utility statement.
2. Provide DD statements with DDNAMEs in the form SW*nn*WK*mm*.
3. Allocate UTPRINT to SYSOUT.

Method 3: Allows the most control over rebuild processing; you must specify both sort work and message data sets.

1. Specify the SORTKEYS keyword in the utility statement.
2. Provide DD statements with DDNAMEs in the form SW*nn*WK*mm*.
3. Provide DD statements with DDNAMEs in the form UTPRIN*nn*.

Data sets used: If you select Method 2 or 3 above, use the information provided here along with “Determining the number of sort subtasks” on page 377, “Allocation of sort subtasks” on page 377, and “Estimating the sort work file size” on page 377 to define the necessary data sets.

Each sort subtask must have its own group of sort work data sets and its own print message data set. Possible reasons to allocate data sets in the utility job JCL rather than using dynamic allocation are:

- To control the size and placement of the data sets.
- To minimize device contention.
- To optimally utilize DASD free space.
- To limit the number of utility subtasks used to build indexes.

The DDNAMEs SW*nn*WK*mm* define the sort work data sets used during utility processing. *nn* identifies the subtask pair, while *mm* identifies one or more data sets to be used by that subtask pair. For example:

SW01WK01 The first sort work data set used by the subtask building the first index.

SW01WK02 The second sort work data set used by the subtask building the first index.

SW02WK01 The first sort work data set used by the subtask building the second index.

SW02WK02 The second sort work data set used by the subtask building the second index.

The DDNAMEs UTPRIN*nn* define the sort work message data sets used by the utility subtask pairs. *nn* identifies the subtask pair.

Determining the number of sort subtasks: The maximum number of utility subtask pairs started for parallel index build is equal to the number of indexes to be built.

REORG TABLESPACE determines the number of subtask pairs according to the following guidelines:

- The number of subtask pairs equals the number of sort work data set groups allocated.
- The number of subtask pairs equals the number of message data sets allocated.
- If you allocate both sort work and message data set groups, the number of subtask pairs equals the smallest number of data sets allocated.

Allocation of sort subtasks: REORG TABLESPACE attempts to assign one sort subtask pair for each index to be built. If REORG TABLESPACE cannot start enough subtasks to build one index per subtask pair, it allocates any excess indexes across the pairs, so one or more subtask pairs might build more than one index.

During parallel index build processing, REORG distributes all indexes among the subtask pairs according to the index creation date, assigning the first created index to the first subtask pair. For SHRLEVEL CHANGE, the mapping index is assigned last.

Estimating the sort work file size: If you choose to provide the data sets, you will need to know the size and number of keys present in all of the indexes being processed by the subtask in order to calculate each sort work file size. After you've determined which indexes are assigned to which subtask pairs, use the following formula to calculate the space required:

2 × (longest index key + 9) × (number of keys extracted)

longest key The length of the longest index key that will be

processed by the subtask. For SHRLEVEL CHANGE, the mapping index key length is 21.

number of keys The number of keys from all indexes to be sorted that will be processed by the subtask.

Keys belonging to partitioning indexes should not be counted in the sort work data set size calculation. The space estimation formula may yield 0 bytes required (because the only index processed by a task set is the partitioning index). In this case, if you are allocating your own sort work data set groups, you still need to allocate sort work data sets for this task set, but you can use a minimal allocation such as 1 track.

Considerations for running REORG

This section discusses additional points to keep in mind when running REORG. See "Appendix C. Resetting an advisory or restrictive status" on page 677 for information on resetting REORG-pending status.

Sorting data in clustering order

When you specify SORTDATA:

- If an explicit clustering index exists on any table in the table space being reorganized, rows of the table space are unloaded in physical sequence. (If the object being reorganized is a partition, rows of that partition are unloaded in physical sequence.) DFSORT then uses the clustering index key to sort the rows. If any other table in the table space has no explicit clustering index, the key of the implicit clustering index (if one exists) is used for the sort.
- If no explicit clustering index exists, SORTDATA is ignored.
- If the largest possible composite record to be sorted exceeds 32760 bytes in length, which is the maximum record size for a BSAM data set, SORTDATA is ignored if you specified SHRLEVEL NONE or REFERENCE (REORG cannot operate if a table has a clustering index and you specified SHRLEVEL CHANGE). This condition can occur only when the table space contains pages of 32 KB size. The largest possible composite record in the table space may be calculated by the following formula:

  ```
  max(K) + max(R + E) + 18 bytes for NONE or REFERENCE or 29 bytes for
  CHANGE
  ```

  where:

  max(K) =
  number of bytes in the longest possible clustering key (K)

  max(R + E) =
  number of bytes in the longest possible record (R) plus edit procedure work area (E). E = 10 bytes if an edit procedure is used, otherwise E = 0 bytes.

When SORTDATA is not specified:

- If an explicit clustering index exists, segmented table spaces are unloaded using that index
- If an explicit clustering index does not exist, the table space is unloaded by table. Multi-table simple table spaces are unloaded by table space scan, in which case rows are reloaded in the same order that they were unloaded.

Methods of unloading data

Data is unloaded by one of three methods:

- *Table space scan with sort*: Chosen if you specified the SORTDATA option, and at least one table in the table space has an explicit clustering index.
- *Table space scan*: Chosen for simple table spaces that contain more than one table, or contain one table but do not have an explicit clustering index.
- *Clustering index*: Always chosen for partitioned table spaces (unless you specified the SORTDATA option); chosen for simple table spaces that contain one table and have an explicit clustering index; and chosen for tables in a segmented table space that have an explicit clustering index.

Encountering an error in the RELOAD phase

Failure during the RELOAD phase (after the data has been unloaded and data sets have been deleted, but before the data has been reloaded) results in an unusable table space.

If the error is on the table space data:

- If you have defined data sets, you can allocate new data sets.
- If STOGROUP has defined data sets, you can alter the new table space to change the primary and secondary quantities.
- If you do allocate new data sets, do alter the table space, or do add volumes to the storage group, restart the REORG job at the beginning of the phase. Otherwise, you can restart either at the last commit point, or at the beginning of the phase.

If the error is on the unloaded data, or if you used the NOSYSREC option, terminate REORG using the TERM UTILITY command. Then recover the table space, using RECOVER, and run the REORG job again.

Reorganizing partitioned table spaces

If you reorganize a single partition or a range of partitions, all indexes of the table space are affected. Depending on how disorganized the nonpartitioning indexes are, you might want to reorganize them as well. For more information about when to reorganize, see “Determining when an index requires reorganization” on page 317.

Reorganizing segmented table spaces

If the target table space is segmented, REORG unloads and reloads by table.

If an explicit clustering index exists on a table in a segmented table space, that table is unloaded in clustering sequence. If NO explicit clustering index exists, the table is unloaded in physical row and segment order.

For segmented table spaces, REORG does NOT normally have to reclaim space from dropped tables. Space freed by dropping tables in a segmented table space is immediately available if the table space can be accessed when DROP TABLE is executed. If the table space cannot be accessed when DROP TABLE is executed, then REORG reclaims the space for dropped tables.

After you run REORG, the segments for each table are contiguous.

Counting records loaded during RELOAD phase

At the end of the RELOAD phase, REORG checks the count of the records that were actually loaded against the count of the records that were unloaded. If the counts do not the match, the actions taken depend on the UNLOAD option you specified on the original job:

- If you specified UNLOAD PAUSE, REORG sets return code 4 and continues processing the job.
- If you specified UNLOAD CONTINUE, DB2 issues an error message and abends the job. The table space or partition remains in RECOVER-pending status.

Reorganizing a LOB table space

Reorganizing a LOB table space is a separate task from reorganizing the base table space. A LOB table space that was defined with LOG YES or LOG NO will affect logging while reorganizing a LOB column. Table 33 on page 204 shows the logging output and LOB table space effect, if any.

Specify LOG YES and SHRLEVEL NONE when you reorganize a LOB table space to avoid leaving the LOB table space in COPY-pending status after the REORG.

Terminating or restarting REORG TABLESPACE

If you terminate REORG TABLESPACE with the TERM UTILITY command during the UNLOAD phase, objects have not yet been changed and the job can be rerun.

If you terminate REORG TABLESPACE with the TERM UTILITY command during the RELOAD phase, the behavior depends on the SHRLEVEL option:

- For SHRLEVEL NONE, the data records are not erased. The table space and indexes are left in RECOVER-pending status. After you recover the table space, rerun the REORG job.
- For SHRLEVEL REFERENCE or CHANGE, the data records are reloaded into shadow objects, so the original objects have not been affected by REORG. You can rerun the job.

If you terminate REORG with the TERM UTILITY command during the SORT, BUILD, or LOG phases, the behavior depends on the SHRLEVEL option:

- For SHRLEVEL NONE, the indexes not yet built are left in RECOVER-pending status. You can run REORG with the SORTDATA option or you can run REBUILD INDEX to rebuild those indexes.
- For SHRLEVEL REFERENCE or CHANGE, the records are reloaded into shadow objects, so the original objects have not been affected by REORG. You can rerun the job.

If you terminate REORG with the TERM UTILITY command during the SWITCH phase, all data sets that were renamed to their shadow counterparts are renamed back, so the objects are left in their original state. You can rerun the job. If there is a problem in renaming to the original data sets, the objects are left in RECOVER-pending status. If the SWITCH phase did not complete, the image copy created by the REORG is not available for use in the recovery. If you terminate an ACTIVE REORG utility with the TERM UTILITY command in the SWITCH phase, during the RENAME process, the RENAME will occur, and the SWITCH phase will complete. The image copy created by the REORG is available for use in future recoveries.

If you terminate REORG with the TERM UTILITY command during the BUILD2 phase, the logical partition is left in RECOVER-pending status. After you run REBUILD INDEX for the logical partition, all objects have completed reorganization.

If you restart a REORG job of certain catalog or directory table spaces, you cannot restart from the last checkpoint. For the following table spaces, you must specify RESTART(PHASE):

DSNDB06.SYSDBASE
DSNDB06.SYSDBAUT
DSNDB06.SYSGROUP
DSNDB06.SYSPLAN
DSNDB06.SYSVIEWS
DSNDB01.DBD01

The REORG-pending status is not reset until the UTILTERM execution phase. If the REORG utility abends or is terminated, the objects are left in REORG pending and RECOVER-pending status, depending on the phase where the failure occurred. See "Appendix C. Resetting an advisory or restrictive status" on page 677 for information about resetting either status.

Table 68. REORG TABLESPACE phases and restrictive statuses

Phase	Effect on restrictive status
UNLOAD	No effect
RELOAD	SHRLEVEL NONE has these effects: • Places table space in RECOVER-pending status at the beginning of the phase and resets the status at the end of the phase. • Places indexes in RECOVER-pending status. • Places the table space in COPY-pending status. If COPYDDN is specified and SORTKEYS is not specified, the COPY-pending status is reset at the end of the phase. SHRLEVEL REFERENCE or CHANGE has no effect.
SORT	No effect
BUILD	SHRLEVEL NONE resets RECOVER-pending status for indexes and, if both COPYDDN and SORTKEYS are specified, resets copy pending status for table spaces at the end of the phase. SHRLEVEL REFERENCE or CHANGE has no effect.
SORTBLD	Same effects as in the SORT and BUILD phases.
LOG	No effect
SWITCH	No effect. Under certain conditions, if TERM UTILITY is issued, it must complete successfully or objects may be placed in RECOVER-pending status.
BUILD2	If TERM UTILITY is issued, the logical partitions for nonpartitioning indexes are placed in logical RECOVER-pending status.

Recovering a failed REORG job: If you terminate REORG SHRLEVEL NONE in the RELOAD phase, all SYSLGRNX records associated with the reorganization are deleted. Use the RECOVER TABLESPACE utility to recover to the current point in time, which recovers the table space to its state before the failed reorganization.

Restarting REORG: Table 69 on page 382 provides information about restarting REORG TABLESPACE, depending on the phase REORG was in when the job stopped.

If you restart REORG in the UTILINIT phase, it re-executes from the beginning of the phase. If REORG abends or system failure occurs while it is in the UTILTERM phase, you must restart the job with RESTART(PHASE).

For each phase of REORG and for each type of REORG TABLESPACE (with SHRLEVEL NONE, with SHRLEVEL REFERENCE, and with SHRLEVEL CHANGE), the table indicates the types of restart allowed (CURRENT and PHASE).

″None″ indicates that no restart is allowed. A blank indicates that the phase does not occur. The ″Data Sets Required″ column lists the data sets that must exist to perform the specified type of restart in the specified phase.

Table 69. REORG TABLESPACE utility restart information

Phase	Type for NONE	Type for REFERENCE	Type for CHANGE	Data Sets Required	Notes
UNLOAD	CURRENT PHASE	CURRENT PHASE	None None	SYSREC	
RELOAD	CURRENT PHASE	CURRENT PHASE	None None	SYSREC and SYSUT1 SYSREC	1,2 1,2
SORT	CURRENT PHASE	CURRENT PHASE	None None	SYSUT1 SYSUT1	2,3 2
BUILD	CURRENT PHASE	CURRENT PHASE	None None	SORTOUT SORTOUT	2,3,4 2,4
SORTBLD	CURRENT PHASE	CURRENT PHASE	None None		2 2
LOG			None None		
SWITCH		CURRENT PHASE	CURRENT PHASE	originals and shadows originals and shadows	3
BUILD2		CURRENT PHASE	CURRENT PHASE	shadows for nonpartitioning indexes shadows for nonpartitioning indexes	3,4 4

Note:

1. For NONE, if you specified NOSYSREC, then RESTART is not possible, and you must execute the RECOVER TABLESPACE utility for the table space or partition. For REFERENCE, if you specified both SORTDATA and NOSYSREC, then RESTART or RESTART(PHASE) restarts at the beginning of the UNLOAD phase.
2. For NONE and REFERENCE, if you specified SORTKEYS, then use RESTART or RESTART(PHASE) to restart at the beginning of the RELOAD phase.
3. You can restart the utility with RESTART or RESTART(PHASE). However, because this phase does not take checkpoints, RESTART restarts from the beginning of the phase.
4. If you specified the PART option with REORG TABLESPACE, you cannot restart the utility at the beginning of the BUILD or BUILD2 phase if any nonpartitioning index is in a page set REBUILD-pending (PSRBD) status.

If you restart a REORG STATISTICS job using RESTART CURRENT, inline statistics collection will not occur. To update catalog statistics, run the RUNSTATS utility after the restarted job completes. Restarting a REORG STATISTICS job with RESTART(PHASE) is conditional after executing UNLOAD PAUSE. To determine if catalog table statistics will be updated using RESTART(PHASE), see Table 70.

Table 70. Statistics collection for REORG TABLESPACE utility phase restart

Phase	CURRENT	PHASE
UTILINIT	NO	YES
UNLOAD	NO	YES
RELOAD	NO	YES
SORT	NO	NO

Table 70. Statistics collection for REORG TABLESPACE utility phase restart (continued)

Phase	CURRENT	PHASE
BUILD	NO	YES
SORTBLD	NO	YES

For instructions on restarting a utility job, see “Chapter 3. Invoking DB2 online utilities” on page 23.

Restarting REORG after an out of space condition: See “Restarting after the output data set is full” on page 47 for guidance in restarting REORG from the last commit point after receiving an out of space condition.

Concurrency and compatibility

Individual data and index partitions, and individual logical partitions of nonpartitioning indexes, are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

REORG of a LOB table space is not compatible with any other utility. The LOB table space is unavailable to other applications during REORG processing.

REORG TABLESPACE compatibility

Table 71 on page 384, Table 72 on page 384, Table 73 on page 385, and Table 74 on page 385 show which claim classes REORG drains and any restrictive state the utility sets on the target object.

For nonpartitioning indexes, REORG PART:
- Drains only the logical partition (and the repeatable read class for the entire index)
- Does not set the page set REBUILD-pending status (PSRCP)
- Does not respect PCTFREE or FREEPAGE attributes when inserting keys

Table 75 on page 386 shows which utilities can run concurrently with REORG on the same target object. The target object can be a table space, an index space, or a partition of a table space or index space. If compatibility depends on particular options of a utility, that is also shown.

Table 76 on page 387 shows which DB2 operations can be affected when reorganizing catalog table spaces.

For SHRLEVEL NONE, Table 71 on page 384 shows which claim classes REORG drains and any restrictive state the utility sets on the target object. For each column, the table indicates the claim or drain that is acquired and the restrictive state that is set in the corresponding phase. UNLOAD CONTINUE and UNLOAD PAUSE, unlike UNLOAD ONLY, include the RELOAD phase and thus include the drains and restrictive states of that phase.

REORG TABLESPACE

Table 71. Claim classes of REORG TABLESPACE SHRLEVEL NONE operations. Use of claims and drains; restrictive states set on the target object.

Target	UNLOAD phase of REORG	RELOAD phase of REORG if UNLOAD CONTINUE or PAUSE	UNLOAD phase of REORG PART	RELOAD phase of REORG PART if UNLOAD CONTINUE or PAUSE
Table space, partition, or a range of partitions of a table space	DW/UTRO	DA/UTUT	DW/UTRO	DA/UTUT
Partitioning index or partition of partitioning index	DW/UTRO	DA/UTUT	DW/UTRO	DA/UTUT
Nonpartitioning index	DW/UTRO	DA/UTUT		DR
Logical partition of nonpartitioning index			DW/UTRO	DA/UTUT

Legend:
- DA: Drain all claim classes, no concurrent SQL access
- DR: Drain the repeatable read class, no concurrent access for SQL repeatable readers
- DW: Drain the write claim class, concurrent access for SQL readers
- UTUT: Utility restrictive state, exclusive control
- UTRO: Utility restrictive state, read only access allowed
- Blank: Any claim, drain, or restrictive state for this object does not change in this phase

For SHRLEVEL REFERENCE, Table 72 shows which claim classes REORG drains and any restrictive state the utility sets on the target object. For each column, the table indicates the claim or drain that is acquired and the restrictive state that is set in the corresponding phase.

Table 72. Claim classes of REORG TABLESPACE SHRLEVEL REFERENCE operations. Use of claims and drains; restrictive states set on the target object.

Target	UNLOAD phase of REORG	SWITCH phase of REORG	UNLOAD phase of REORG PART	SWITCH phase of REORG PART	BUILD2 phase of REORG PART
Table space or partition of table space	DW/UTRO	DA/UTUT	DW/UTRO	DA/UTUT	UTRW
Partitioning index or partition of partitioning index	DW/UTRO	DA/UTUT	DW/UTRO	DA/UTUT	UTRW
Nonpartitioning index	DW/UTRO	DA/UTUT		DR	
Logical partition of nonpartitioning index			DW/UTRO	DA/UTUT	

Legend:
- DA: Drain all claim classes, no concurrent SQL access
- DDR: Dedrain the read claim class, concurrent SQL access
- DR: Drain the repeatable read class, no concurrent access for SQL repeatable readers
- DW: Drain the write claim class, concurrent access for SQL readers
- UTUT: Utility restrictive state, exclusive control
- UTRO: Utility restrictive state, read only access allowed
- Blank: Any claim, drain, or restrictive state for this object does not change in this phase

For REORG of an entire table space with SHRLEVEL CHANGE, Table 73 shows which claim classes REORG drains and any restrictive state the utility sets on the target object.

Table 73. Claim classes of REORG TABLESPACE SHRLEVEL CHANGE operations. Use of claims and drains; restrictive states set on the target object.

Target	UNLOAD phase	Last iteration of LOG phase	SWITCH phase
Table space	CR/UTRW[1]	DW/UTRO	DA/UTUT
Index	CR/UTRW[1]	DW/UTRO	DA/UTUT

Legend:
- CR: Claim the read claim class
- DA: Claim all claim classes, no concurrent SQL access
- DW: Drain the write claim class, concurrent access for SQL readers
- UTUT: Utility restrictive state, exclusive control
- UTRO: Utility restrictive state, read only access allowed
- UTRW: Utility restrictive state, read/write access allowed

Notes:

1. If the target object is a segmented table space, SHRLEVEL CHANGE does not allow you to concurrently execute an SQL *searched* DELETE without the WHERE clause.

For REORG of a partition with SHRLEVEL NONE, Table 74 shows which claim classes REORG drains and any restrictive state the utility sets on the target object.

Table 74. Claim classes of REORG TABLESPACE SHRLEVEL CHANGE operations. Use of claims and drains; restrictive states set on the target object.

Target	UNLOAD phase	Last iteration of LOG phase	SWITCH phase	BUILD2 phase
Partition of table space	CR/UTRW	DW/UTRO or DA/UTUT[1]	DA/UTUT	UTRW
Partition of partitioning index	CR/UTRW	DW/UTRO or DA/UTUT[1]	DA/UTUT	UTRW
Nonpartitioning index			DR	
Logical partition of nonpartitioning index	CR/UTRW	DW/UTRO or DA/UTUT[1]	DA/UTUT	

Legend:
- CR: Claim the read claim class
- DA: Drain all claim classes, no concurrent SQL access
- DDR: Dedrain the read claim class, no concurrent access for SQL repeatable readers
- DR: Drain the repeatable read class, no concurrent access for SQL repeatable readers
- DW: Drain the write claim class, concurrent access for SQL readers
- UTUT: Utility restrictive state, exclusive control
- UTRO: Utility restrictive state, read only access allowed
- UTRW: Utility restrictive state, read/write access allowed
- Blank: Any claim, drain, or restrictive state for this object does not change in this phase

Notes:

1. DA/UTUT applies if you specified DRAIN ALL.

REORG TABLESPACE

Table 75. REORG TABLESPACE compatibility

Action	REORG SHRLEVEL NONE UNLOAD CONTINUE or PAUSE, REORG SHRLEVEL REFERENCE, or REORG SHRLEVEL CHANGE	REORG SHRLEVEL NONE UNLOAD ONLY without cluster index	REORG SHRLEVEL NONE UNLOAD ONLY with cluster index
CATMAINT	No	No	No
CHECK DATA	No	No	No
CHECK INDEX	No	Yes	Yes
CHECK LOB	No	No	No
COPY INDEXSPACE	No	Yes	Yes
COPY TABLESPACE	No	Yes	Yes
DIAGNOSE	Yes	Yes	Yes
LOAD	No	No	No
MERGECOPY	No	No	No
MODIFY	No	No	No
QUIESCE	No	Yes	Yes
REBUILD INDEX	No	Yes	No
RECOVER INDEX	No	Yes	No
RECOVER INDEXSPACE	No	No	No
RECOVER TABLESPACE	No	No	No
REORG INDEX	No	Yes	No
REORG TABLESPACE SHRLEVEL NONE UNLOAD CONTINUE or PAUSE, REORG SHRLEVEL REFERENCE, or REORG SHRLEVEL CHANGE	No	No	No
REORG TABLESPACE SHRLEVEL NONE UNLOAD ONLY or EXTERNAL	No	Yes	Yes
REPAIR DUMP or VERIFY	No	Yes	Yes
REPAIR LOCATE KEY or RID DELETE or REPLACE	No	No	No
REPAIR LOCATE TABLESPACE PAGE REPLACE	No	No	No
REPAIR LOCATE INDEX PAGE REPLACE	No	Yes	No
REPORT	Yes	Yes	Yes
RUNSTATS	No	Yes	Yes
STOSPACE	No	Yes	Yes
UNLOAD	No	Yes	Yes

Table 76. DB2 operations affected by reorganizing catalog table spaces

Catalog Table Space	Actions That Might Not Run Concurrently
Any table space except SYSCOPY and SYSSTR	CREATE, ALTER, and DROP statements
SYSCOPY, SYSDBASE,SYSDBAUT, SYSSTATS, SYSUSER	Utilities
SYSDBASE, SYSDBAUT, SYSGPAUT, SYSPKAGE, SYSPLAN, SYSUSER	GRANT and REVOKE statements
SYSDBAUT, SYSDBASE, SYSGPAUT, SYSPKAGE, SYSPLAN, SYSSTATS, SYSUSER, SYSVIEWS	BIND and FREE commands
SYSPKAGE, SYSPLAN	Plan or package execution

Reviewing REORG TABLESPACE output

The output from REORG TABLESPACE consists of a reorganized table space, partition, or a range of partitions; from REORG INDEX it consists of a reorganized index or index partition. Table 77 summarizes the effect of REORG on a table space partition and on the corresponding index partition.

Table 77. REORG summary

Specification	Results
REORG TABLESPACE	All data + entire partitioning index + all nonpartitioning indexes
REORG TABLESPACE PART *n*	Data for part *n* + part *n* of the partitioning index + index entries for part *n* in all nonpartitioning indexes
REORG TABLESPACE PART *n1*:*n2*	Data for parts *n1* through *n2* + parts *n1* through *n2* of the partitioned index + index entries for those parts in all nonpartitioning indexes

When reorganizing a segmented table space, REORG leaves free pages and free space on each page in accordance with the current values of the FREEPAGE and PCTFREE parameters. (Those values can be set by the CREATE TABLESPACE, ALTER TABLESPACE, CREATE INDEX, or ALTER INDEX statements). REORG leaves one free page after reaching the FREEPAGE limit for each table in the table space. When reorganizing a nonsegmented table space, REORG leaves one free page after reaching the FREEPAGE limit, regardless of whether the records loaded belong to the same or different tables.

Segments that contain a table that has an explicit cluster index are unloaded using the cluster index; when the table is loaded, all data records are in cluster key order.

After running REORG TABLESPACE

After a reorganization has completed:

- If you have used LOG YES, consider taking an image copy of the reorganized table space or partition to:
  - Provide a full image copy for recovery. This prevents having to process the log records written during reorganization.
  - Permit making incremental image copies later.

You might not need to take an image copy of a table space for which all the following are true:
- It is relatively small
- It is used only in read-only applications
- It can be easily loaded again in the event of failure.

See “Chapter 9. COPY” on page 81 for information on making image copies.

- If you use REORG SHRLEVEL NONE LOG NO on a LOB table space and the LOB manager determines that nothing needs to be done to the table space, no COPY-pending status is set. However, if the LOB manager indicates that changes are needed, REORG places the reorganized LOB table space or partition in COPY-pending status. In this situation, perform a full image copy to reset the COPY-pending status and to ensure that a backup is available for recovery.

  You should also run the COPY utility if the REORG was performed to turn off REORG-pending status (REORP), and an inline copy was not taken. You can not use an image copy created before turning off REORP.
- If you use COPYDDN, SHRLEVEL REFERENCE, or SHRLEVEL CHANGE, and the object you are reorganizing is not a catalog or directory table space for which COPYDDN is ignored, you do not need to take an image copy.
- Use the RUNSTATS utility on the table space and its indexes if inline statistics were not collected, so that the DB2 catalog statistics take into account the newly reorganized data, and SQL paths can be selected with accurate information. You only need to run RUNSTATS on nonpartitioning indexes if you reorganized a subset of the partitions.
- If you have used REORG TABLESPACE SHRLEVEL CHANGE, you can drop the mapping table and its index.
- If you have used SHRLEVEL REFERENCE or CHANGE, and a table space, partition, or index resides in user-managed data sets, you can delete the user-managed shadow data sets.
- If you specified DISCARD on a REORG of a table involved in a referential integrity set, you will need to run CHECK DATA against any affected referentially-related objects that were placed in CHECK-pending status.

Sample control statements

Example 1: REORG using default sort output data set. This example shows the DDNAME for the unload data set is UNLD, the DDNAME for the sort input data set is WORK, and the DDNAME for the sort output data set is defaulted to SORTOUT.

```
//STEP1    EXEC DSNUPROC,UID='IUJLU101.REORG',
//         UTPROC='',
//         SYSTEM='V71A'
//UTPRINT  DD  SYSOUT=*
//UNLD     DD DSN=IUJLU101.REORG.STEP1.UNLD,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTWK01  DD DSN=IUJLU101.REORG.STEP1.SORTWK01,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTWK02  DD DSN=IUJLU101.REORG.STEP1.SORTWK02,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//WORK   DD DSN=IUJLU101.REORG.STEP1.WORK,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD DSN=IUJLU101.REORG.STEP1.SORTOUT,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
REORG TABLESPACE (DSN8D71A.DSN8S71D)
```

```
      SORTDATA
      UNLDDN (UNLD)
      WORKDDN (WORK)
//*
```

Example 2: Reorganizing a table space. Reorganize table space DSN8S71D in database DSN8D71A.

```
REORG TABLESPACE DSN8D71A.DSN8S71D
         SORTDATA
```

Example 3: Reorganizing a table space partition. Reorganize partition 3 of table space DSN8S71E in database DSN8D71A.

```
REORG TABLESPACE DSN8D71A.DSN8S71E
  PART 3
  SORTDATA
  SORTDEVT SYSDA
```

Example 4: REORG unloading by table space scan and then SORT. Reorganize table space DSN8S71E in database DSN8D71A. Specify that the data is unloaded by table space scan and then sorted before being reloaded.

```
REORG TABLESPACE DSN8D71A.DSN8S71E SORTDATA
```

Example 5: REORG Using SORTKEYS. Use the SORTKEYS option to invoke parallel index build for a reorganization of the table space DSN8S71D in database DSNDB04. This example does not specify that dynamic allocation is to be used by DFSORT. Instead, it allocates sort work data sets in two groups, which limits the number of utility subtask pairs to two. This example does not require UTPRIN*nn* DD statements, because it uses DSNUPROC to invoke utility processing, which includes a DD statement that allocates UTPRINT to SYSOUT.

```
//SAMPJOB  JOB  ...
//STEP1    EXEC DSNUPROC,UID='SAMPJOB.REORG',UTPROC='',SYSTEM='V71A'
//SYSREC   DD DSN=SAMPJOB.REORG.STEP1.SYSREC,DISP=(NEW,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//* First group of sort work data sets for parallel index build
//SW01WK01 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SW01WK02 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SW01WK03 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//* Second group of sort work data sets for parallel index build
//SW02WK01 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SW02WK02 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SW02WK03 DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//* Sort work data sets for use by SORTDATA
//SORTWK01 DD  UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SORTWK02 DD  UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SORTWK03 DD  UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND)
//SYSIN    DD *
REORG TABLESPACE DSNDB04.DSN8S71D LOG NO SORTDATA SORTKEYS
/*
```

Example 6: REORG Using SORTKEYS while allowing read-write access. Use the SORTKEYS option to invoke parallel index build for reorganization of the table space DSN8S71E in database DSNDB04 in database DSNDB04. The name of the mapping table is DSN8MAP. DFSORT dynamically allocates sort work data sets. This example does not require UTPRIN*nn* DD statements, because it uses DSNUPROC to invoke utility processing, which includes a DD statement that allocates UTPRINT to SYSOUT.

```
//SAMPJOB  JOB  ...
//STEP1    EXEC DSNUPROC,UID='SAMPJOB.REORG',UTPROC='',SYSTEM='DSN'
//SYSCOPY  DD UNIT=SYSDA,SPACE=(CYL,(10,20),,,ROUND),
//            DSN=SAMPJOB,COPY,DISP=(NEW,CATLG,CATLG)
```

```
//SYSIN    DD *
REORG TABLESPACE DSNDB04.DSN8S71E LOG NO SORTDEVT SYSDA SORTNUM 4
  SHRLEVEL CHANGE MAPPINGTABLE DSN8MAP
/*
```

Example 7: Reorganizing a table while allowing read-only access. Reorganize table space DSN8S71D in database DSN8D71A. The deadline for start of the SWITCH phase is *3:15* on *February 4, 1997.*

```
REORG TABLESPACE DSN8D71A.DSN8S71D COPYDDN(MYCOPY1)
  RECOVERYDDN(MYCOPY2) SHRLEVEL REFERENCE
  DEADLINE 1997-2-4-03.15.00
```

Example 8: Reorganizing a table while allowing read-write access. Reorganize table space DSN8S71D in database DSN8D71A. The deadline for start of the SWITCH phase is *3:15* on February 4, 1997. The name of the mapping table is MYMAPTABLE. The maximum desired amount of time for the log processing in the read-only (last) iteration of log processing is *240* seconds. If reorganization's reading of the log is not catching up to applications' writing of the log quickly enough, DB2 will drain the write claim class after sending the LONGLOG message to the operator. That draining will take place at least 900 seconds after the LONGLOG message is sent.

```
REORG TABLESPACE DSN8D71A.DSN8S71D COPYDDN(MYCOPY1)
  RECOVERYDDN(MYCOPY2) SHRLEVEL CHANGE
  DEADLINE 1997-2-4-03.15.00
  MAPPINGTABLE
DSN8710.MAP_TBL MAXRO 240 LONGLOG DRAIN DELAY 900
```

Example 9: Reorganizing a range of table space partitions. Reorganize partitions 3 through 5 of table space DSN8S71E in database DSN8D71A.

```
REORG TABLESPACE DSN8D71A.DSN8S71E
  PART 3:5
  STATISTICS
  SORTDEVT SYSDA
    SHRLEVEL NONE
    COPYDDN SYSCOPY
    SORTDATA
```

Example 10: REORG a partition using STATISTICS. Reorganize partition 3 of table space DSN8S71E in database DSN8D71A, using the STATISTICS option to update catalog table statistics for that table.

```
REORG TABLESPACE DSN8D71A.DSN8S71E
   SORTDATA STATISTICS PART 3
```

Example 11: REORG using STATISTICS to update table space and index statistics. Reorganize table space DSN8S71E in database DSN8D71A, using the STATISTICS option to update catalog statistics for the table space and all indexes defined on that table.

```
REORG TABLESPACE DSN8D71A.DSN8S71E SORTDATA STATISTICS
   TABLE
   INDEX(ALL) KEYCARD FREQVAL NUMCOLS 1
   COUNT 10 REPORT YES UPDATE NONE
```

Example 12: Checking if a table should be reorganized. Report if the OFFPOSLIMIT or INDREFLIMIT values are exceeded for the TPHR5201 table space in database DBHR5201.

```
//STEP1    EXEC DSNUPROC,UID='HUHRU252.REORG2',TIME=1440,
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//SYSREC   DD DSN=HUHRU252.REORG2.STEP1.SYSREC,DISP=(MOD,DELETE,CATLG),
```

```
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSCOPY  DD DSN=HUHRU252.REORG2.STEP1.SYSCOPY,DISP=(MOD,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD DSN=HUHRU252.REORG2.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD DSN=HUHRU252.REORG2.STEP1.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
REORG TABLESPACE DBHR5201.TPHR5201 PART 11
      SORTDATA SORTKEYS NOSYSREC
      REPORTONLY
      SHRLEVEL CHANGE MAPPINGTABLE ADMF001.MAP1
      COPYDDN (SYSCOPY)
          OFFPOSLIMIT 11 INDREFLIMIT 15
/*
```

On successful completion, DB2 returns output similar to the following output:

```
DSNU050I   DSNUGUTC - REORG TABLESPACE DBHR5201.TPHR5201 PART 11 SORTDATA SORTKEYS NOSYSREC REPORTONLY SHRLEVEL CHA
NGE MAPPINGTABLE ADMF001.MAP1 COPYDDN(SYSCOPY) OFFPOSLIMIT 11 INDREFLIMIT 15 STATISTICS UPDATE SPACE
DSNU286I ( DSNURLIM - REORG TABLESPACE DBHR5201.TPHR5201 OFFPOSLIMIT SYSINDEXPART ROWS
  CREATOR .IXNAME             CREATOR .TBNAME             PART CARDF       FAROFFPOSF  NEAROFFPOSF      STATSTIME
* ADMF001 .IPHR5201           ADMF001 .TBHR5201             11  6.758E+03   2.972E+03   7.38E+02   1999-02-05-08.27.04
DSNU287I ( DSNURLIM - REORG TABLESPACE DBHR5201.TPHR5201 INDREFLIMIT SYSTABLEPART ROWS
  DBNAME  .TSNAME   PART         CARD  FARINDREF NEARINDREF      STATSTIME
  DBHR5201.TPHR5201   11         6758          0          0 1999-02-05-08.27.04
DSNU289I ( DSNURLIM -  REORG LIMITS HAVE BEEN MET
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=2
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = HUHRU252.REORG2
DSNU050I   DSNUGUTC - REORG INDEX ADMF001.IPHR5201 PART 11 LEAFDISTLIMIT 2 REPORTONLY
DSNU288I ( DSNURLIM - REORG INDEX ADMF001.IPHR5201 LEAFDISTLIMIT SYSINDEXPART ROWS
  CREATOR .IXNAME             PART   LEAFDIST      STATSTIME
* ADMF001 .IPHR5201             11          3 1999-02-05-08.27.04
DSNU289I ( DSNURLIM -  REORG LIMITS HAVE BEEN MET
DSNU050I   DSNUGUTC - CHECK INDEX(ADMF001.IPHR5201 PART 11)
DSNU700I ( DSNUKGET - 6761 INDEX ENTRIES UNLOADED FROM INDEX='ADMF001.IPHR5201' PARTITION= 11
DSNU705I   DSNUK001 - UNLOAD PHASE COMPLETE - ELAPSED TIME=00:00:00
DSNU042I   DSNUGSOR - SORT PHASE STATISTICS -
                      NUMBER OF RECORDS=6761
                      ELAPSED TIME=00:00:01
DSNU717I ( DSNUKTER - 6761 ENTRIES CHECKED FOR INDEX 'ADMF001.IPHR5201' PARTITION= 11
DSNU720I   DSNUK001 - CHECKIDX PHASE COMPLETE, ELAPSED TIME=00:00:02
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=2
```

Figure 21. Sample output showing REORG limits have been met

Example 13: Conditionally reorganizing a table. To ensure recent statistics for the table space, execute the RUNSTATS utility for the TPHR5201 table space. Then, reorganize the TPHR5201 table space in database DBHR5201 if the OFFPOSLIMIT or INDREFLIMIT value is exceeded.

```
//*******************************************************************
//* COMMENT: UPDATE STATISTICS
//*******************************************************************
//STEP1    EXEC DSNUPROC,UID='HUHRU252.REORG1',TIME=1440,
//         UTPROC='',
//         SYSTEM=V71AR',DB2LEV=DB2A
//SYSREC   DD DSN=HUHRU252.REORG1.STEP1.SYSREC,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD DSN=HUHRU252.REORG1.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD DSN=HUHRU252.REORG1.STEP1.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
RUNSTATS TABLESPACE DBHR5201.TPHR5201
      STATISTICS UPDATE SPACE
```

```
/*
//*****************************************************************
//* COMMENT: REORG THE TABLESPACE
//*****************************************************************
//STEP2    EXEC DSNUPROC,UID='HUHRU252.REORG1',TIME=1440,
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//SYSREC   DD DSN=HUHRU252.REORG1.STEP1.SYSREC,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSCOPY1 DD DSN=HUHRU252.REORG1.STEP1.SYSCOPY1,
//         DISP=(MOD,CATLG,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//SYSUT1   DD DSN=HUHRU252.REORG1.STEP1.SYSUT1,DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD DSN=HUHRU252.REORG1.STEP1.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
//         SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
REORG TABLESPACE DBHR5201.TPHR5201
      SORTDATA NOSYSREC SORTKEYS
      COPYDDN SYSCOPY1
      OFFPOSLIMIT
      INDREFLIMIT
      STATISTICS TABLE(ALL) INDEX(ALL)
/*
```

On successful completion, DB2 returns output similar to the following output:

```
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = HUHRU252.REORG1
DSNU050I   DSNUGUTC - RUNSTATS TABLESPACE DBHR5201.TPHR5201 UPDATE SPACE
DSNU610I - DSNUSUTP -   SYSTABLEPART CATALOG UPDATE FOR DBHR5201.TPHR5201 SUCCESSFUL
DSNU620I - DSNUSDRA - RUNSTATS CATALOG TIMESTAMP = 1999-05-03-14.57.58.921242
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = HUHRU252.REORG1
DSNU050I   DSNUGUTC - REORG TABLESPACE DBHR5201.TPHR5201 SORTDATA NOSYSREC SORTKEYS COPYDDN(SYSCOPY1) OFFPOSLIMIT
INDREFLIMIT STATISTICS TABLE ALL INDEX(ALL)
DSNU286I - DSNURLIM - REORG TABLESPACE DBHR5201.TPHR5201 OFFPOSLIMIT SYSINDEXPART ROWS
  CREATOR .IXNAME   CREATOR .TBNAME        PART CARDF       FAROFFPOSF  NEAROFFPOSF      STATSTIME
* ADMF001 .IPHR5201    ADMF001 .TBHR5201        1  3.6E+01     8.0E+00     4.0E+00    1999-05-03-14.57.43
  ADMF001 .IPHR5201    ADMF001 .TBHR5201        2  5.0E+00     0.0E0       0.0E0      1999-05-03-14.57.43
  ADMF001 .IPHR5201    ADMF001 .TBHR5201        3  5.4E+01     0.0E0       0.0E0      1999-05-03-14.57.43
  ADMF001 .IPHR5201    ADMF001 .TBHR5201        4  3.0E+01     1.0E+00     0.0E0      1999-05-03-14.57.43
* ADMF001 .IPHR5201    ADMF001 .TBHR5201        5  2.1E+01     2.0E+00     1.0E+00    1999-05-03-14.57.43
  ADMF001 .IPHR5201    ADMF001 .TBHR5201        6  5.0E+00     0.0E0       0.0E0      1999-05-03-14.57.43
  ADMF001 .IPHR5201    ADMF001 .TBHR5201        7  4.0E+00     0.0E0       0.0E0      1999-05-03-14.57.43
* ADMF001 .IPHR5201    ADMF001 .TBHR5201        8  3.5E+01     9.0E+00     8.0E+00    1999-05-03-14.57.43
* ADMF001 .IPHR5201    ADMF001 .TBHR5201        9  2.5E+01     4.0E+00     1.0E+00    1999-05-03-14.57.43
  ADMF001 .IPHR5201    ADMF001 .TBHR5201       10  1.0E+00     0.0E0       0.0E0      1999-05-03-14.57.43
* ADMF001 .IPHR5201    ADMF001 .TBHR5201       11  6.758E+03   2.972E+03   7.38E+02   1999-05-03-14.57.43
DSNU287I - DSNURLIM - REORG TABLESPACE DBHR5201.TPHR5201 INDREFLIMIT SYSTABLEPART ROWS
  DBNAME  .TSNAME   PART         CARD  FARINDREF NEARINDREF      STATSTIME
  DBHR5201.TPHR5201    1           36          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201    2            5          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201    3           54          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201    4           30          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201    5           21          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201    6            5          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201    7            4          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201    8           35          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201    9           25          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201   10            1          0          0 1999-05-03-14.57.58
  DBHR5201.TPHR5201   11         6758          0          0 1999-05-03-14.57.58
DSNU289I - DSNURLIM - REORG LIMITS HAVE BEEN MET
DSNU290I - DSNURLIM - REORG WILL BE PERFORMED
DSNU252I   DSNUGSRT - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS UNLOADED=6974 FOR TABLESPACE DBHR5201.TPHR5201
DSNU250I   DSNUGSRT - UNLOAD PHASE COMPLETE, ELAPSED TIME=00:00:03
DSNU400I   DSNURBID - COPY PROCESSED FOR TABLESPACE DBHR5201.TPHR5201
                      NUMBER OF PAGES=1034
                      AVERAGE PERCENT FREE SPACE PER PAGE = 14.96
                      PERCENT OF CHANGED PAGES =100.00
                      ELAPSED TIME=00:01:46
DSNU610I - DSNUSUTP - SYSTABLEPART CATALOG UPDATE FOR DBHR5201.TPHR5201 SUCCESSFUL
DSNU610I - DSNUSUPT - SYSTABSTATS CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUPC - SYSCOLSTATS CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUTB - SYSTABLES CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUCO - SYSCOLUMNS CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUTS - SYSTABLESPACE CATALOG UPDATE FOR DBHR5201.TPHR5201 SUCCESSFUL
DSNU620I - DSNURDRT - RUNSTATS CATALOG TIMESTAMP = 1999-05-03-14.58.16.924784
DSNU304I - DSNURWT -  (RE)LOAD PHASE STATISTICS - NUMBER OF RECORDS=6974 FOR TABLE ADMF001.TBHR5201
DSNU302I   DSNURILD - (RE)LOAD PHASE STATISTICS - NUMBER OF INPUT RECORDS PROCESSED=6974
DSNU300I   DSNURILD - (RE)LOAD PHASE COMPLETE, ELAPSED TIME=00:01:48
DSNU042I   DSNUGSOR - SORT PHASE STATISTICS -
                      NUMBER OF RECORDS=27896
                      ELAPSED TIME=00:00:00
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=36 FOR INDEX ADMF001.IPHR5201  PART 1
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=5 FOR INDEX ADMF001.IPHR5201  PART 2
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=54 FOR INDEX ADMF001.IPHR5201  PART 3
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=30 FOR INDEX ADMF001.IPHR5201  PART 4
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=21 FOR INDEX ADMF001.IPHR5201  PART 5
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=5 FOR INDEX ADMF001.IPHR5201  PART 6
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=4 FOR INDEX ADMF001.IPHR5201  PART 7
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=35 FOR INDEX ADMF001.IPHR5201  PART 8
```

Figure 22. Sample showing RUNSTATS and conditional REORG output (Part 1 of 2)

```
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=25 FOR INDEX ADMF001.IPHR5201  PART 9
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=1 FOR INDEX ADMF001.IPHR5201  PART 10
DSNU348I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=6758 FOR INDEX ADMF001.IPHR5201  PART 11
DSNU349I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=6974 FOR INDEX ADMF001.IXHR521B
DSNU349I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=6974 FOR INDEX ADMF001.IXHR521C
DSNU349I - DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=6974 FOR INDEX ADMF001.IXHR521D
DSNU610I - DSNUSUIP - SYSINDEXPART CATALOG UPDATE FOR ADMF001.IPHR5201 SUCCESSFUL
DSNU610I - DSNUSUPI - SYSINDEXSTATS CATALOG UPDATE FOR ADMF001.IPHR5201 SUCCESSFUL
DSNU610I - DSNUSUPD - SYSCOLDISTSTATS CATALOG UPDATE FOR ADMF001.IPHR5201 SUCCESSFUL
DSNU610I - DSNUSUPC - SYSCOLSTATS CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUIX - SYSINDEXES CATALOG UPDATE FOR ADMF001.IPHR5201 SUCCESSFUL
DSNU610I - DSNUSUCO - SYSCOLUMNS CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUCD - SYSCOLDIST CATALOG UPDATE FOR ADMF001.IPHR5201 SUCCESSFUL
DSNU610I - DSNUSUIP - SYSINDEXPART CATALOG UPDATE FOR ADMF001.IXHR521B SUCCESSFUL
DSNU610I - DSNUSUIX - SYSINDEXES CATALOG UPDATE FOR ADMF001.IXHR521B SUCCESSFUL
DSNU610I - DSNUSUCO - SYSCOLUMNS CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUCD - SYSCOLDIST CATALOG UPDATE FOR ADMF001.IXHR521B SUCCESSFUL
DSNU610I - DSNUSUIP - SYSINDEXPART CATALOG UPDATE FOR ADMF001.IXHR521C SUCCESSFUL
DSNU610I - DSNUSUIX - SYSINDEXES CATALOG UPDATE FOR ADMF001.IXHR521C SUCCESSFUL
DSNU610I - DSNUSUCO - SYSCOLUMNS CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUCD - SYSCOLDIST CATALOG UPDATE FOR ADMF001.IXHR521C SUCCESSFUL
DSNU610I - DSNUSUIP - SYSINDEXPART CATALOG UPDATE FOR ADMF001.IXHR521D SUCCESSFUL
DSNU610I - DSNUSUIX - SYSINDEXES CATALOG UPDATE FOR ADMF001.IXHR521D SUCCESSFUL
DSNU610I - DSNUSUCO - SYSCOLUMNS CATALOG UPDATE FOR ADMF001.TBHR5201 SUCCESSFUL
DSNU610I - DSNUSUCD - SYSCOLDIST CATALOG UPDATE FOR ADMF001.IXHR521D SUCCESSFUL
DSNU620I - DSNURDRI - RUNSTATS CATALOG TIMESTAMP = 1999-05-03-15.00.05.010574
DSNU258I   DSNURBXD - BUILD PHASE STATISTICS - NUMBER OF INDEXES=4
DSNU259I   DSNURBXD - BUILD PHASE COMPLETE, ELAPSED TIME=00:00:11
DSNU405I   DSNURORG - DB2 IMAGE COPY SUCCESSFUL FOR TABLESPACE DBHR5201.TPHR5201
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=2
```

Figure 22. Sample showing RUNSTATS and conditional REORG output (Part 2 of 2)

Example 14: Reorganizing a table space using DRAIN WAIT, RETRY and RETRY_DELAY. Use the DRAIN WAIT option to specify that the utility will wait for 30 seconds for DML. Use RETRY to indicate that the utility will retry four times and RETRY_DELAY to specify that the interval between retries is 10 seconds.

```
//STEP1    EXEC DSNUPROC,UID='HUHRU257.REORG',TIME=1440,
//         UTPROC='',
//         SYSTEM='SSTR',DB2LEV=DB2A
//UTPRINT  DD  SYSOUT=*
//SYSIN    DD *
   OPTIONS PREVIEW
   TEMPLATE CPYTMP UNIT(SYSDA)
                   DSN(HUHRU257.REORG.T&TI..SYSCOPY1)
   TEMPLATE SREC
            UNIT(SYSDA) DISP(NEW,CATLG,CATLG)
            DSN(HUHRU257.REORG.&ST..SREC)
   TEMPLATE SDISC
            UNIT(SYSDA) DISP(NEW,CATLG,CATLG)
            DSN(HUHRU257.REORG.&ST..SDISC)
   TEMPLATE SPUNCH
            UNIT(SYSDA) DISP(NEW,CATLG,CATLG)
            DSN(HUHRU257.REORG.&ST..SPUNCH)
   TEMPLATE SUT1
            UNIT(SYSDA) DISP(NEW,DELETE,CATLG)
            DSN(HUHRU257.REORG.&ST..SUT1)
   TEMPLATE SOUT
            UNIT(SYSDA) DISP(NEW,DELETE,CATLG)
            DSN(HUHRU257.REORG.&ST..SOUT)
   LISTDEF REORG_TBSP INCLUDE TABLESPACE DBHR5701.TPHR5701
   OPTIONS OFF
   REORG TABLESPACE LIST REORG_TBSP
       DRAIN_WAIT 30    RETRY 4     RETRY_DELAY 10
```

```
STATISTICS
TABLE (ALL) SAMPLE 60
INDEX (ALL KEYCARD FREQVAL NUMCOLS 2 COUNT 15)
SHRLEVEL CHANGE MAPPINGTABLE MAP5702
LONGLOG CONTINUE MAXRO 400   DELAY 30
COPYDDN (CPYTMP)
SORTKEYS SORTDEVT SYSDA SORTNUM 8
PUNCHDDN SPUNCH
DISCARDDN SDISC
UNLDDN SREC
WORKDDN (SUT1,SOUT)
```

Chapter 24. REPAIR

The REPAIR online utility repairs data. The data can be your own data, or data you would not normally access, such as space map pages and index entries.

REPAIR is intended as a means of replacing invalid data with valid data. Be extremely careful using REPAIR. Improper use can damage the data even further.

You can use the REPAIR utility to:
- Test DBDs
- Repair DBDs
- Reset a pending status on a table space or index
- Verify the contents of data areas in table spaces and indexes
- Replace the contents of data areas in table spaces and indexes
- Delete a single row from a table space
- Produce a hexadecimal dump of an area in a table space or index
- Delete an entire LOB from a LOB table space
- Dump LOB pages
- Rebuild OBDs for a LOB table space

For a diagram of REPAIR syntax and a description of available options, see “Syntax and options of the control statement” on page 398. For detailed guidance on running this utility, see “Instructions for running REPAIR” on page 409.

Output: The potential output from the REPAIR utility consists of a modified page or pages in the specified DB2 table space or index, and a dump of the contents.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- REPAIR privilege for the database
- DBADM or DBCTRL authority for the database
- SYSCTRL or SYSADM authority.

An ID with installation SYSOPR authority can also execute REPAIR, but only on a table space in the DSNDB01 or DSNDB06 database.

To execute REPAIR DBD, the privilege set must include SYSADM, SYSCTRL, or installation SYSOPR authority.

REPAIR should only be used by a knowledgeable person. Be careful to grant REPAIR authorization only to appropriate people.

Execution phases of REPAIR: One of the following phases can be identified if the job terminates.

The phases for REPAIR are:

Phase	Description
UTILINIT	Initialization
REPAIR	
UTILTERM	Cleanup.

© Copyright IBM Corp. 1983, 2001

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

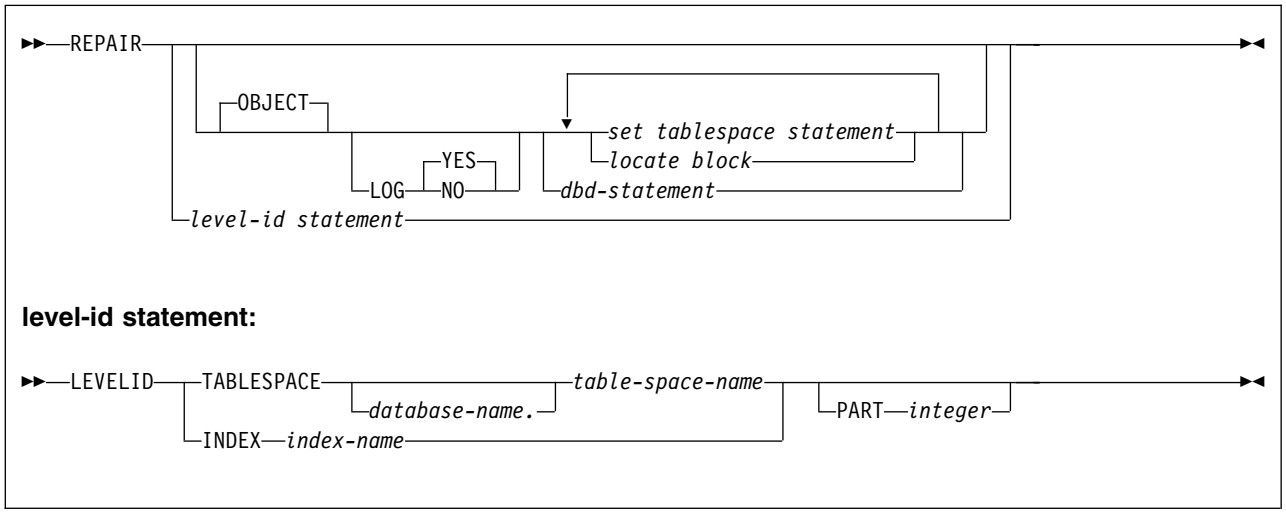

REPAIR option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

OBJECT
Is optional, and used for clarity only.

LOG
Tells whether to log the changes made by REPAIR. If the changes are logged, they are applied again if the data is recovered.

YES
Logs the changes.

The **default** is **YES**.

REPAIR LOG YES cannot override the LOG NO attribute of a table space.

NO
Does not log the changes. You cannot use this option with a DELETE statement.

REPAIR LOG NO can override the LOG YES attribute of a table space.

LEVELID
Sets the level identifier of the named table space, table space partition, index, or index space partition to a new identifier. Use LEVELID to accept the use of a down-level data set. You cannot specify multiple LEVELIDs.

You cannot use LEVELID with a table space, table space partition, index, or index space partition with outstanding indoubt log records or pages in the logical page list (LPL).

Attention: Accepting the use of a down-level data set might cause data inconsistencies. Problems with inconsistent data resulting from resetting the level identifier are the responsibility of the user.

TABLESPACE *database-name.table-space-name*
Specifies the table space (and, optionally, the database to which it belongs) whose level identifier is set.

database-name
Specifies the name of the database the table space belongs to.

The **default** is **DSNDB04**.

table-space-name
Specifies the name of the table space.

INDEX
Specifies the index whose level identifier is set.

index-name
Specifies the index to be processed.

PART
Identifies a partition of the table space or index.

integer is the number of the partition and must be in the range from one to the number of partitions defined for the object. The maximum is 254.

SET TABLESPACE and SET INDEX statement syntax

The SET TABLESPACE statement resets the COPY-pending, RECOVER-pending, CHECK-pending, auxiliary warning (AUXW), and auxiliary CHECK-pending (ACHKP) statuses for a table space or data set. The SET INDEX statement resets the informational COPY-pending (ICOPY), RECOVER-pending, REBUILD-pending, or CHECK-pending status for an index.

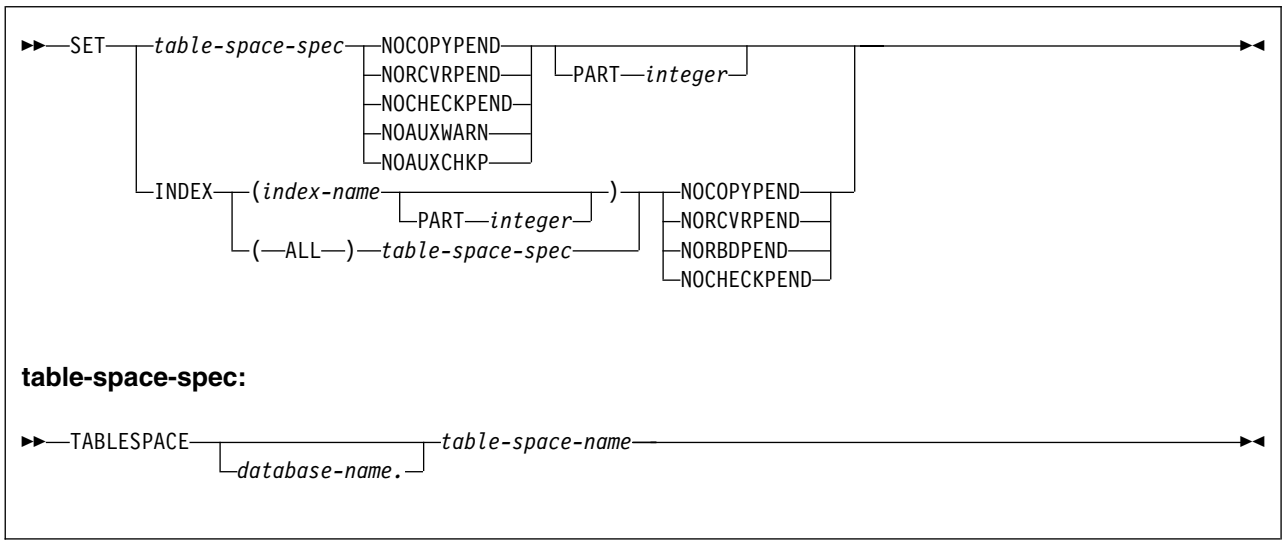

SET TABLESPACE and SET INDEX option descriptions

TABLESPACE *database-name.table-space-name*
Specifies the table space (and, optionally, the database to which it belongs) whose pending status is to be reset.

database-name
Specifies the name of the database the table space belongs to.

The **default** is **DSNDB04**.

table-space-name
Specifies the name of the table space.

INDEX
Specifies the index whose RECOVER-pending, CHECK-pending, REBUILD-pending, or informational COPY-pending status is to be reset.

(*index-name***)**
Specifies the index to be processed.

(ALL) Specifies that all indexes in the table space will be processed.

The keyword INDEXES is still accepted following a *table-space-spec* and causes all indexes to be processed. All indexes can also be processed by specifying INDEX(ALL) followed by a *table-space-spec*.

NOCOPYPEND
Resets the COPY-pending status of the specified table space, or informational COPY-pending (ICOPY) status of the specified index.

NORCVRPEND
Resets the RECOVER-pending (RECP) status of the specified table space or index.

NORBDPEND
Resets the REBUILD-pending (RBDP) status, the page set REBUILD-pending status (PSRBDP), or the RBDP* status of the specified index.

NOCHECKPEND
Resets the CHECK-pending (CHKP) status of the specified table space or index.

NOAUXWARN
Resets the auxiliary warning (AUXW) status of the specified table space. The specified table space must be a base table space or a LOB table space.

NOAUXCHKP
Resets the auxiliary CHECK-pending (ACHKP) status of the specified table space. The specified table space must be a base table space.

PART *integer*
Specifies a particular partition whose COPY-pending, informational COPY-pending, or RECOVER-pending status is to be reset. If you do not specify PART, REPAIR resets the pending status of the entire table space or index.

integer is the number of the partition and must be in the range from one to the number of partitions defined for the object. The maximum is 254.

You can specify PART for NOCHECKPEND on a table space, and for NORCVRPEND on indexes.

The PART keyword is not valid for a LOB table space or an index on the auxiliary table.

LOCATE block syntax

A LOCATE block is a set of statements, each with its own options, that begins with a LOCATE statement and ends with the next LOCATE or SET statement, or with the end of the job. There can be more than one LOCATE block in a REPAIR utility statement.

In any LOCATE block, you can use VERIFY, REPLACE, or DUMP as often as you like; you can use DELETE only once.

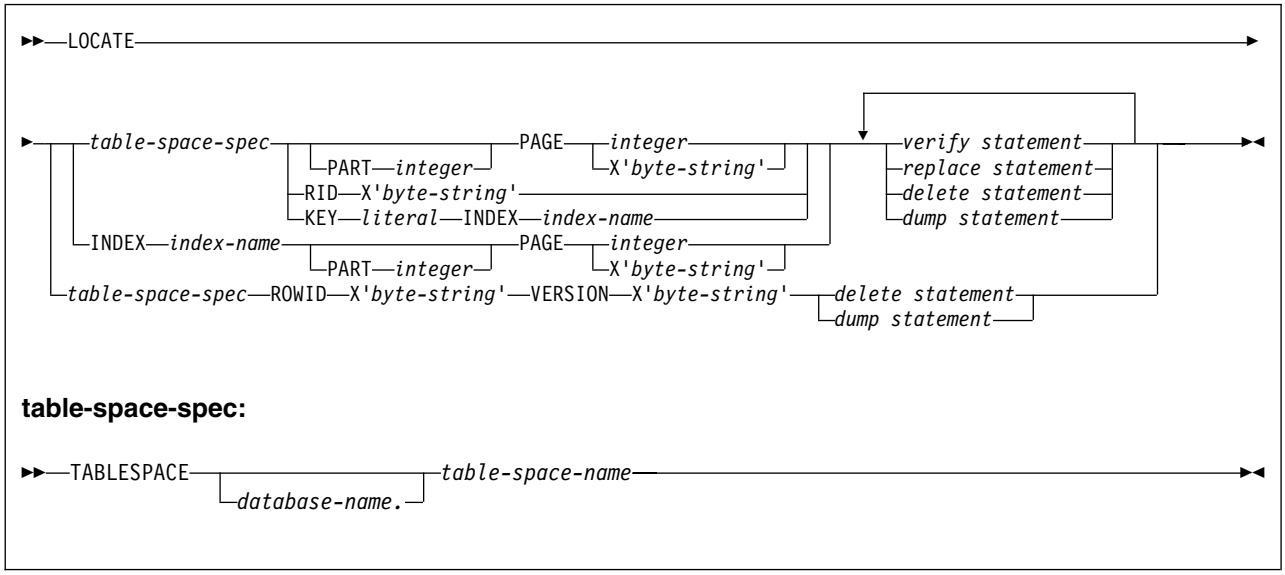

LOCATE TABLESPACE statement option descriptions

The LOCATE TABLESPACE statement locates data to be repaired within a table space.

One LOCATE statement is required for each unit of data to be repaired. Several LOCATE statements can appear after each REPAIR statement.

If a REPAIR statement is followed by several LOCATE statements, all processing caused by VERIFY, REPLACE, and DUMP statements is committed before the next LOCATE is processed.

TABLESPACE *database-name.table-space-name*
Specifies the table space (and, optionally, the database to which it belongs) in which data is to be located for repair.

database-name
Is the name of the database to which the table space belongs and is optional.

The **default** is **DSNDB04**.

table-space-name
Is the name of the table space containing the data you want to repair.

PART *integer*
Valid only for partitioned table spaces.

integer is the number of the partition containing the page to be located.

PAGE
Specifies the relative page number within the table space, partitioned table space, or index that is to be operated on. The first page in either case is 0 (zero).

integer
integer is a decimal number from one to six digits in length.

X'*byte-string*'
Specifies that the data of interest is an entire page. The offsets given in *byte-string* and in subsequent statements are relative to the beginning of the page. The first byte of the page is 0.

byte-string is a value from one to eight hexadecimal characters in length. You do not need to enter leading zeros. Enclose the *byte-string* between apostrophes and precede it with X.

RID X'*byte-string*'
Specifies that the data to be located is a single row. The offsets given in *byte-string* and in subsequent statements are relative to the beginning of the row. The first byte of the stored row prefix is 0.

byte-string can be one to eight hexadecimal characters. You do not need to enter leading zeros. Enclose the byte string between single quotation marks and precede it with an X.

KEY *literal*
Specifies that the data to be located is a single row, identified by *literal*. The offsets given in subsequent statements are relative to the beginning of the row. The first byte of the stored row prefix is at offset 0.

literal is any SQL constant that can be compared with the key values of the named index.

Character constants specified within the LOCATE KEY option may not be specified as American National Standard Code for Information Interchange (ASCII) or UNICODE character strings. No conversion of the values is performed. To use this option when the table space is ASCII or UNICODE, you should specify the values as hexadecimal constants.

If more than one row has the value *literal* in the key column, REPAIR returns a list of record identifiers (RIDs) for records with that key value, but does NOT perform any other operations (verify, replace, delete, or dump) until the next LOCATE TABLESPACE statement is encountered. To repair the proper data, write a LOCATE TABLESPACE statement that selects the desired row, using the RID option, the PAGE option, or a different KEY and INDEX option. Then execute REPAIR again.

INDEX *index-name*
Specifies a particular index that is to be used to find the row containing the key. When you are locating by key, the index you specify must be single-column.

index-name is the qualified or unqualified name of the index.

ROWID X'*byte-string*'
Specifies that the data to be located is a LOB in a LOB table space.

byte-string is the row ID that identifies the LOB column.

The most likely source of a row ID is an orphaned LOB that is reported by the CHECK LOB utility. If you specify the ROWID keyword, the specified table space must be a LOB table space.

VERSION X'*byte-string*'
Specifies that the data to be located is a LOB in a LOB table space.

byte-string is the version number that identifies the version of the LOB column.

The most likely source of a version number is an orphaned LOB that is reported by the CHECK LOB utility, or an out-of-synch LOB that is reported by the CHECK DATA utility. If you specify the VERSION keyword, the specified table space must be a LOB table space.

LOCATE INDEX statement option descriptions

The LOCATE INDEX statement locates data to be repaired within an index.

One LOCATE statement is required for each unit of data to be repaired. Several LOCATE statements can appear after each REPAIR statement.

If a REPAIR statement is followed by several LOCATE statements, all processing caused by VERIFY, REPLACE, and DUMP statements is committed before the next LOCATE is processed.

index-name
Specifies the index containing the data you want to repair.

index-name is the qualified name of the index, in the form *creator-id.index-name*. If you omit the qualifier creator ID, the user identifier for the utility job is used.

PART *integer*
Valid only for indexes of partitioned table spaces.

integer is the number of the partitioning index containing the page to be located.

PAGE *integer*
Specifies the relative page number within the index space to be operated on. The first page is 0 (zero).

integer
integer is a decimal number from one to six digits in length.

X'*byte-string*'
Specifies that the data of interest is an entire page. The offsets given in *byte-string* and in subsequent statements are relative to the beginning of the page. The first byte of the page is 0.

byte-string is a value from one to eight hexadecimal characters in length. You do not need to enter leading zeros. Enclose the *byte-string* between apostrophes and precede it with X.

VERIFY statement syntax

The VERIFY statement tests whether a particular data area contains a specified value. If the data area *does* contain the value, subsequent operations in the same

LOCATE block are allowed to proceed. If *any* data area *does not* contain its specified value, *all* subsequent operations in the same LOCATE block are inhibited.

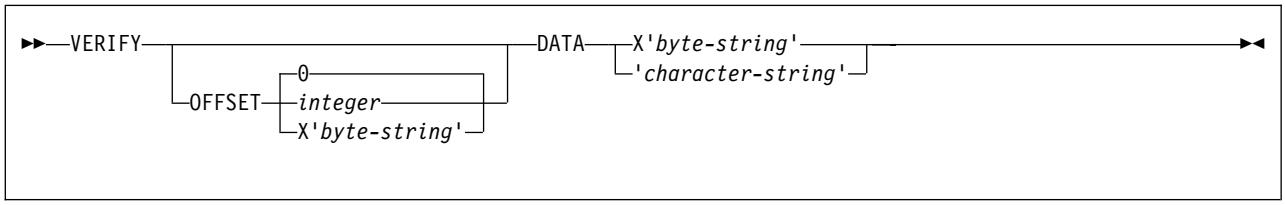

VERIFY statement option descriptions

OFFSET
Locates the data to be tested by a relative byte address within the row or page.

integer
Gives the offset as an integer. The **default** is **0**, the first byte of the area identified by the previous LOCATE statement.

X'*byte-string***'**
Gives the offset as one to four hexadecimal characters. You need not enter leading zeros. Enclose the byte string between single quotation marks and precede it with X.

DATA
Tells what data is assumed to be present before a change is made.

Character constants specified within the VERIFY DATA option may not be specified as ASCII or UNICODE character strings. No conversion of the values is performed. To use this option when the table space is ASCII or UNICODE, you should specify the values as hexadecimal constants.

X'*byte-string***'**
Can be an *even number*, from two to thirty-two, of hexadecimal characters that must be present. You need not enter leading zeros. Enclose the byte string between single quotation marks and precede it with X.

'*character-string***'**
Can be any character string that must be present.

REPLACE statement syntax

The REPLACE statement replaces data at a particular location. The statement falls within a LOCATE block. If any VERIFY statement within that block finds a data area that does not contain its specified data, the REPLACE operation is inhibited.

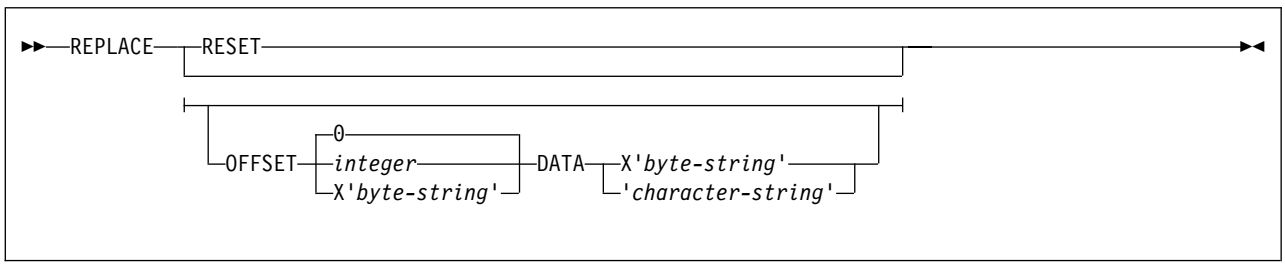

REPLACE statement option descriptions

RESET
Resets the inconsistent data indicator. A page for which this indicator is on is considered in error, and the indicator must be reset before you can access the page. Numbers of pages with inconsistent data are reported at the time they are encountered.

The option also resets the PGCOMB flag bit in the first byte of the page to agree with the bit code in the last byte of the page.

OFFSET
Tells where data is to be replaced by a relative byte address within the row or page.

integer
Gives the offset as an integer. The **default** is **0**, the first byte of the area identified by the previous LOCATE statement.

X'*byte-string*'
Gives the offset as one to four hexadecimal characters. You need not enter leading zeros. Enclose the byte string between single quotation marks and precede it with X.

DATA
Defines the new data that is to be entered. Only one OFFSET and one DATA specification are acted on for each REPLACE statement.

Character constants specified within the VERIFY DATA option may not be specified as ASCII or UNICODE character strings. No conversion of the values is performed. To use this option when the table space is ASCII or UNICODE, you should specify the values as hexadecimal constants.

X'*byte-string*'
Can be an *even number*, from two to thirty-two, of hexadecimal characters to replace the current data. You need not enter leading zeros. Enclose the byte string between single quotation marks and precede it with X.

'*character-string*'
Can be any character string to replace the current data.

DELETE statement syntax and description

The DELETE statement deletes a single row of data that has been located by a RID or KEY option. The statement falls within a LOCATE block. If any VERIFY statement within that block finds a data area that does not contain its specified data, the DELETE operation is inhibited.

The DELETE statement operates without regard for referential constraints. If you delete a parent row, its dependent rows remain unchanged in the table space. However, in the DB2 catalog and directory table spaces, where links are used to reference other tables in the catalog, deleting a parent row causes all child rows to be deleted as well. Moreover, deleting a child row in the DB2 catalog tables also updates its predecessor and successor pointer to reflect the deletion of this row. Therefore, if the child row has incorrect pointers, the DELETE might lead to an unexpected result. See page 416 for a possible method of deleting a child row without updating its predecessor and successor.

In any LOCATE block, you can use DELETE only once.

You cannot use DELETE if you have used any of these options:

- The LOG NO option on the REPAIR statement
- A LOCATE INDEX statement to begin the LOCATE block
- The PAGE option on the LOCATE TABLESPACE statement in the same LOCATE block
- A REPLACE statement for the same row of data.

When you specify LOCATE ROWID for a LOB table space, the LOB specified by ROWID is deleted with its index entry. All pages occupied by the LOB are converted to free space. The DELETE statement will *not* remove any reference to the deleted LOB from the base table space.

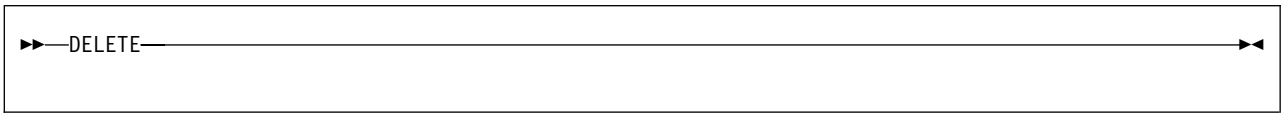

DUMP statement syntax

The DUMP statement produces a hexadecimal dump of data identified by offset and length. DUMP statements have no effect on VERIFY or REPLACE operations.

When you specify LOCATE ROWID for a LOB table space, one or more map or data pages of the LOB are dumped. The DUMP statement dumps all of the LOB column pages if you do not specify either the MAP or DATA keyword.

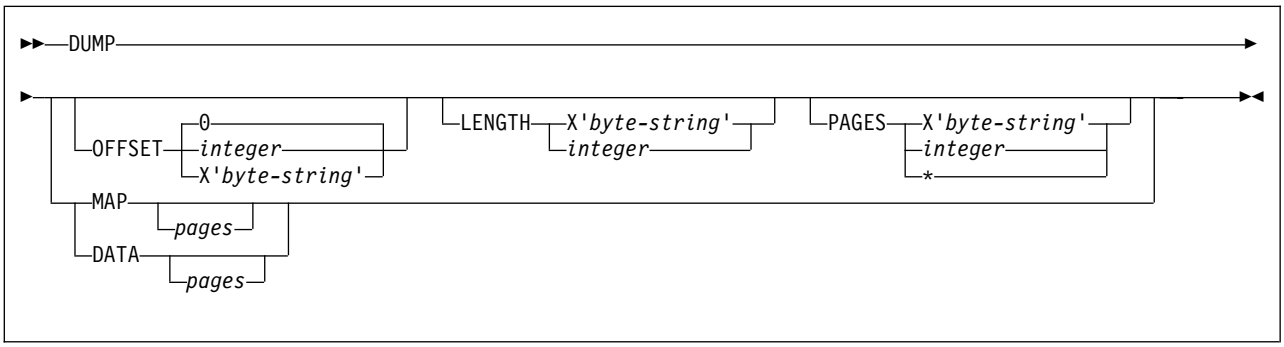

DUMP statement option descriptions

OFFSET
Optionally, locates the data to be dumped by a relative byte address within the row or page.

integer
Gives the offset as an integer. The **default** is **0**, the first byte of the row or page.

X'*byte-string***'**
Gives the offset as one to four hexadecimal characters. You need not enter leading zeros. Enclose the byte string between single quotation marks and precede it with X.

LENGTH
Optionally, tells the number of bytes of data to dump. If you omit both LENGTH and PAGE, the dump continues from OFFSET to the end of the row or page.

If you give a number of bytes (with LENGTH) and a number of pages (with PAGE) the dump contains the same relative bytes from each page. That is, from each page you see the same number of bytes, at the same offset.

X'*byte-string***'**
Can be one to four hexadecimal characters. You need not enter leading zeros. Enclose the byte string between single quotation marks and precede it with X.

integer
Gives the number as an integer.

PAGES
Optionally, tells a number of pages to dump. You can use this option only if you used PAGE in the preceding LOCATE TABLESPACE control statement.

X'*byte-string***'**
Can be one to four hexadecimal characters. You need not enter leading zeros. Enclose the byte string between single quotation marks and precede it with X.

integer
Gives the number as an integer.

*****
Dumps all pages from the starting point to the end of the table space or partition.

MAP *pages*
Dumps only the LOB map pages.

pages specifies the number of LOB map pages to dump. If you do not specify *pages*, all LOB map pages of the LOB that is specified by ROWID and version are dumped.

DATA *pages*
Dumps only the LOB data pages.

pages specifies the number of LOB data pages to dump. If you do not specify *pages*, all LOB data pages of the LOB that is specified by ROWID and version are dumped.

DBD statement syntax

The DBD statement allows you to:

- Compare the definition of a database in the DB2 catalog with its definition in the DB2 directory
- Rebuild a database definition in the directory from the information in the DB2 catalog, including LOB information
- Drop an inconsistent database definition from the DB2 catalog and the DB2 directory.

REPAIR also assumes that the links in table spaces DSNDB01.DBD01, DSNDB06.SYSDBAUT, DSNDB06.SYSDBASE are intact. Before executing REPAIR with the DBD statement, run the DSN1CHKR utility (page 555) on these table spaces to ensure that the links are not broken.

The database on which REPAIR DBD is run must be started for access by utilities only. For more information about using the DBD statement, see page “Using the DBD statement” on page 411.

You can use REPAIR DBD on declared temporary tables, which must be created in a database that is defined with the AS TEMP clause. No other DB2 utilities can be used on a declared temporary table, its indexes, or its table spaces.

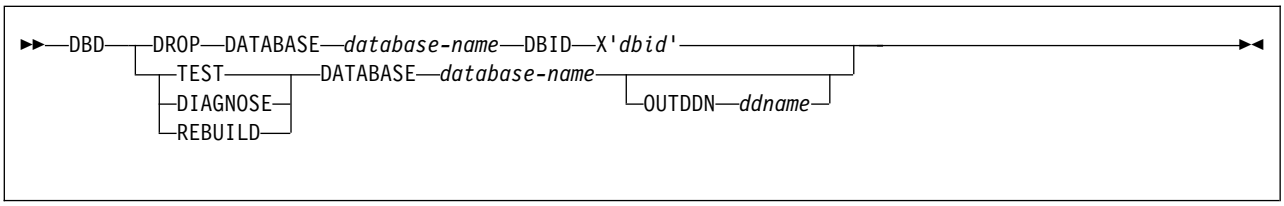

DBD statement option descriptions

DROP
Drops the specified database from both the DB2 catalog and the DB2 directory. Use this keyword if the SQL DROP DATABASE statement fails because the description of the database is not in both the DB2 catalog and the DB2 directory. If you cannot use the ALTER command to remove the WITH RESTRICT ON DROP option on tables in a database that is badly damaged and you need to drop the database, you can use this keyword to drop the database.

Attention: Use the DROP option with extreme care. For further assistance, you can contact the IBM Support Center.

DATABASE *database-name*
Specifies the target database.

database-name is the name of the target database, which cannot be DSNDB01 (the DB2 directory) or DSNDB06 (the DB2 catalog).

If you use TEST, DIAGNOSE, or REBUILD, *database-name* cannot be DSNDB07 (the work file database).

If you use DROP, *database-name* cannot be DSNDB04 (the default database).

DBID X'*dbid***'**
Specifies the database descriptor identifier for the target database.

dbid is the database descriptor identifier.

TEST
Builds a DBD from information in the DB2 catalog, and compares it with the DBD in the DB2 directory. TEST reports significant differences between the two DBDs.

If the condition code is 0, then the DBD in the DB2 directory is consistent with the information in the DB2 catalog.

If the condition code is not 0, then the information in the DB2 catalog and the DBD in the DB2 directory can be inconsistent. Run REPAIR DBD with the DIAGNOSE option to gather information necessary for resolving any possible inconsistency.

DIAGNOSE
Produces information necessary for resolving an inconsistent database definition. Like the TEST option, DIAGNOSE builds a DBD based on the information in the DB2 catalog and compares it with the DBD in the DB2 directory. In addition, DIAGNOSE reports any differences between the two DBDs, and produces hexadecimal dumps of the inconsistent DBDs.

If the condition code is 0, then the information in the DB2 catalog and the DBD in the DB2 directory is consistent.

If the condition code is 8, then the information in the DB2 catalog and the DBD in the DB2 directory can be inconsistent.

Use your electronic link with IBM Support Center, if available, for help in resolving any inconsistencies.

REBUILD
Rebuilds the DBD associated with the specified database from the information in the DB2 catalog.

Attention: Use the REBUILD option with extreme care. For further assistance, you can contact the IBM Support Center.

OUTDDN *ddname*
Specifies the DD statement for an optional output data set. This data set contains copies of the DB2 catalog records used to rebuild the DBD.

ddname is the name of the DD statement.

Instructions for running REPAIR

To run REPAIR, you must:

1. Read “Before running REPAIR” in this chapter.
2. Prepare the necessary data sets, as described in 397.
3. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for REPAIR, see “Sample control statements” on page 416.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 410. (For a complete description of the syntax and options for REPAIR, see “Syntax and options of the control statement” on page 398.)
5. Check the compatibility table in “Concurrency and compatibility” on page 413 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the REPAIR job doesn’t complete, as described in “Terminating or restarting REPAIR” on page 413.
7. Run REPAIR.

Attention: Be extremely careful when using the REPAIR utility to replace data. Changing data to invalid values using REPLACE might produce unpredictable results, particularly when changing page header information. Improper use of REPAIR can result in damaged data, or in some cases, system failure.

See “Chapter 3. Invoking DB2 online utilities” on page 23 for an explanation of ways to execute DB2 utilities.

Before running REPAIR

Making a copy of the table space

Before starting to use REPAIR to change data, it may be useful to have a copy (full image copy or DSN1COPY) of the affected table space to enable fallback.

Restoring damaged indexes

Because REPAIR can only access index data by referring to a page and an offset within the page, identifying a problem and correcting it can be difficult. Use REBUILD INDEX or RECOVER INDEX to restore damaged index data.

Data sets used by REPAIR

Table 78 describes the data sets used by REPAIR. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 78. Data sets used by REPAIR

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
Optional output data set	Data set containing copies of the DB2 catalog records use to rebuild the DBD. The DD name is defined by the user.	No

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table space
Object to be repaired. It is named in the REPAIR control statement and is accessed through the DB2 catalog.

Calculating output data set size: Use the following calculation to estimate the size of the output data set:

SPACE = (4096,(n,n))

where *n* = total number of records in your catalog relating to the database on which REPAIR DBD is being executed.

An estimate for *n* can be calculated by summing the results of SELECT COUNT(*) from all of the catalog tables in the SYSDBASE table space where the name of the database associated with the record matches the database on which REPAIR DBD is being executed.

Creating the control statement

See “Syntax and options of the control statement” on page 398 for REPAIR syntax and option descriptions. See “Sample control statements” on page 416 for examples of REPAIR usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

“Resetting table space status” on page 411
“Repairing a damaged page” on page 411
“Using the DBD statement” on page 411
“Locating rows by key” on page 412
“Using VERIFY, REPLACE, and DELETE operations” on page 412
“Repairing critical catalog table spaces and indexes” on page 412

Resetting table space status

In most cases, it is better to reset the COPY-pending restriction by taking a full image copy rather than using REPAIR. This is because RECOVER can not be executed successfully until an image copy has been made.

It is better to reset the RECOVER-pending status by running RECOVER or LOAD rather than using REPAIR. This is because RECOVER uses DB2 controlled recovery information, while REPAIR SET TABLESPACE or INDEX resets the recovery pending status without considering the recoverability of the table space, such as the availability of image copies, of rows in SYSIBM.SYSCOPY, and of log data sets.

It is better to verify and possibly correct referential integrity constraints by running CHECK DATA. CHECK DATA performs a complete check of all referential integrity constraints of the table space set, while REPAIR leaves you with the responsibility of checking all the referential integrity constraints violations.

To reset the CHECK-pending status for a LOB table space, it is recommended that you:

1. Run the CHECK DATA utility again with the AUXERROR INVALIDATE keywords specified, then
2. Update the invalid LOBs.

To reset the auxiliary warning (AUXW) status for a LOB table space, it is recommended that you:

1. Update the invalid LOBs, then
2. Run the CHECK LOB utility.

Resetting index space status

It is better to run COPY INDEXSPACE to reset the informational COPY-pending status than to use the REPAIR utility to reset the status.

Consider using the REBUILD INDEX or RECOVER INDEX utility on an index that is in REBUILD-pending status, rather than running REPAIR SET INDEX NORBDPEND. This is because RECOVER uses DB2 controlled recovery information, while REPAIR SET INDEX resets the rebuild-pending status without considering the recoverability of the index, such as the availability of image copies, of rows in SYSIBM.SYSCOPY, and of log data sets.

Repairing a damaged page

1. Execute REPAIR with the LOG YES option and the DUMP control statement, specifying the pages you suspect are damaged. Verify that the dump you receive contains the desired pages.
2. If you know which page is damaged and you can see how to resolve the error, repair the page and reset the “inconsistent data” indicator. Run REPAIR with the REPLACE RESET DATA control statement. Keep track of your actions in case you need to undo anything later.
3. If you determined that the page is not *really* damaged, but merely has the “inconsistent data” indicator on, reset the indicator by running REPAIR with the REPLACE RESET control statement.

Using the DBD statement

The following is the recommended procedure for using the DBD statement:

1. Run the DSN1CHKR utility on the DSNDB01.DBD01, DSNDB06.SYSDBAUT, and DSNDB06.SYSDBASE table spaces.

2. Issue the STOP DATABASE (*database-name*) command, then issue the START DATABASE (*database-name*) ACCESS(UT) command to allow only utilities to access the database that is associated with the DBD.
3. Run REPAIR DBD with the TEST option to determine if the information in the DB2 catalog is consistent with the DBD in the DB2 directory.
4. If inconsistencies exist (condition code is not 0), use the DIAGNOSE option with the OUTDDN keyword to produce diagnostic information. Use your electronic link with IBM Support Center, if available, for assistance in analyzing this information.
5. You might be instructed to replace the existing DBD with the REBUILD option. DO NOT use this option if you suspect that information in the catalog is causing the inconsistency. REBUILD uses information in the catalog to rebuild the DBD; if the catalog is incorrect, the rebuilt DBD will be incorrect.

   DB2 reads each table space in the database during the REBUILD process to gather information. If the data sets for the table spaces do not exist or are not accessible to DB2, then the REBUILD abends.
6. If you think there is an inconsistency in the DBD of the work file database, run REPAIR DBD DROP or DROP DATABASE (SQL) and then recreate it. If you receive errors when you drop the work file database, contact your IBM Support Center for assistance.

Locating rows by key

If you use LOCATE TABLESPACE KEY, a number of rows might satisfy the condition. In this case, REPAIR only returns the RIDs of the rows, and does not perform any VERIFY, REPLACE, DELETE or DUMP which might be coded in that LOCATE block. The RID option of LOCATE TABLESPACE can then be used to identify a specific row. Examples of the messages issued are shown below:

```
DSNU658I - DSNUCBRL - MULTIPLE RECORDS FOUND WITH SPECIFIED KEY
DSNU660I - DSNUCBRL - POSSIBLE RID - X00000100B'
DSNU660I - DSNUCBRL - POSSIBLE RID - X000000C18'
DSNU660I - DSNUCBRL - POSSIBLE RID - X000000916'
DSNU660I - DSNUCBRL - POSSIBLE RID - X000000513'
DSNU650I - DSNUCBRP - DUMP
DSNU012I   DSNUGBAC - UTILITY EXECUTION TERMINATED,
                      HIGHEST RETURN CODE=8
```

Multiple column indexes: The KEY option only supports single column indexes. The following message will be received if an attempt is made to locate a row using a multiple column index.

```
DSNUCBRK - INDEX USED HAS MULTIPLE-FIELD KEY
```

Using VERIFY, REPLACE, and DELETE operations

If *any* data area *does not* contain the value required by a VERIFY statement, *all* REPLACE and DELETE operations in the same locate block are inhibited. VERIFY and REPLACE statements following the next LOCATE are not affected.

Repairing critical catalog table spaces and indexes

An ID with a granted authority receives message DSNT500I, “RESOURCE UNAVAILABLE”, while trying to repair a table space or index in the catalog or directory if table space DSNDB06.SYSDBASE or DSNDB06.SYSUSER is unavailable. If you get this message, you must either make these table spaces available or run the REPAIR utility on the catalog or directory using an authorization ID with the installation SYSADM or installation SYSOPR authority.

Terminating or restarting REPAIR

REPAIR can be terminated with the TERM UTILITY command. See Chapter 2 of *DB2 Command Reference* for information about TERM UTILITY.

REPAIR cannot be restarted. If you attempt to restart REPAIR, you receive message DSNU181I, and the job abends. You must terminate the job with the TERM UTILITY command, and rerun REPAIR from the beginning.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Table 79 shows which claim classes REPAIR drains and any restrictive state the utility sets on the target object.

Table 80 on page 414 and Table 81 on page 414 show which utilities can run concurrently with REPAIR on the same target object. The target object can be a table space, an index space, or a partition of a table space or index space. If compatibility depends on particular options of a utility, that is also shown.

Table 79. Claim classes of REPAIR operations. Use of claims and drains; restrictive states set on the target object.

Action	Table space or partition	Index or partition
REPAIR LOCATE KEY DUMP or VERIFY	DW/UTRO	DW/UTRO
REPAIR LOCATE KEY DELETE or REPLACE	DA/UTUT	DA/UTUT
REPAIR LOCATE RID DUMP or VERIFY	DW/UTRO	
REPAIR LOCATE RID DELETE	DA/UTUT	DA/UTUT
REPAIR LOCATE RID REPLACE	DA/UTUT	
REPAIR LOCATE TABLESPACE DUMP or VERIFY	DW/UTRO	
REPAIR LOCATE TABLESPACE REPLACE	DA/UTUT	
REPAIR LOCATE INDEX PAGE DUMP or VERIFY		DW/UTRO
REPAIR LOCATE INDEX PAGE DELETE		DA/UTUT

Legend:
- DA - Drain all claim classes - no concurrent SQL access
- DW - Drain the write claim class - concurrent access for SQL readers
- UTUT - Utility restrictive state - exclusive control
- UTRO - Utility restrictive state - read only access allowed
- Blank - Object is not affected by this utility.

REPAIR does not set a utility restrictive state if the target object is DSNDB01.SYSUTILX.

Table 80. Compatibility with REPAIR, LOCATE by KEY or RID

	DUMP or VERIFY	DELETE or REPLACE
CHECK DATA	No	No
CHECK INDEX	Yes	No
CHECK LOB	Yes	No
COPY INDEXSPACE	Yes	No
COPY TABLESPACE	Yes	No
DIAGNOSE	Yes	Yes
LOAD	No	No
MERGECOPY	Yes	Yes
MODIFY	Yes	Yes
QUIESCE	Yes	No
REBUILD INDEX	No	No
RECOVER INDEX	No	No
Note: RECOVER INDEX is compatible with LOCATE by RID, DUMP or VERIFY.		
RECOVER TABLESPACE	No	No
REORG INDEX	No	No
Note: REORG INDEX is compatible with LOCATE by RID, DUMP, VERIFY, or REPLACE.		
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No	No
REORG TABLESPACE UNLOAD ONLY or EXTERNAL	Yes	No
REPAIR DUMP or VERIFY	Yes	No
REPAIR DELETE or REPLACE	No	No
Note: REPAIR LOCATE INDEX PAGE REPLACE is compatible with LOCATE by RID REPLACE.		
REPORT	Yes	Yes
RUNSTATS INDEX SHRLEVEL REFERENCE	Yes	No
RUNSTATS INDEX SHRLEVEL CHANGE	Yes	Yes
RUNSTATS TABLESPACE	Yes	No
STOSPACE	Yes	Yes
UNLOAD	Yes	NO

Table 81. Compatibility with REPAIR, LOCATE by PAGE

	TABLESPACE DUMP or VERIFY	TABLESPACE REPLACE	INDEX DUMP or VERIFY	INDEX REPLACE
SQL Read	Yes	No	Yes	No
SQL Write	No	No	No	No
CHECK DATA	No	No	No	No

Table 81. Compatibility with REPAIR, LOCATE by PAGE (continued)

	TABLESPACE DUMP or VERIFY	TABLESPACE REPLACE	INDEX DUMP or VERIFY	INDEX REPLACE
CHECK INDEX	Yes	No	Yes	No
CHECK LOB	Yes	No	Yes	No
COPY INDEXSPACE	Yes	Yes	Yes	No
COPY TABLESPACE	Yes	No	Yes	No
DIAGNOSE	Yes	Yes	Yes	Yes
LOAD	No	No	No	No
MERGECOPY	Yes	Yes	Yes	Yes
MODIFY	Yes	Yes	Yes	Yes
QUIESCE	Yes	No	Yes	No
REBUILD INDEX	Yes	No	No	n/a
RECOVER INDEX	Yes	No	No	No
RECOVER TABLESPACE (no option)	No	No	Yes	Yes
RECOVER TABLESPACE TOCOPY or TORBA	No	No	No	No
RECOVER TABLESPACE ERROR RANGE	No	No	Yes	Yes
REORG INDEX	Yes	Yes	No	No
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No	No	No	No
REORG TABLESPACE UNLOAD ONLY or EXTERNAL	Yes	No	Yes	Yes
REPAIR DUMP or VERIFY	Yes	No	Yes	No
REPAIR DELETE or REPLACE	No	No	No	No
Note: REPAIR LOCATE INDEX PAGE REPLACE is compatible with LOCATE TABLESPACE PAGE.				
REPORT	Yes	Yes	Yes	Yes
RUNSTATS INDEX	Yes	Yes	Yes	No
RUNSTATS TABLESPACE	Yes	No	Yes	Yes
STOSPACE	Yes	Yes	Yes	Yes
UNLOAD	Yes	No	Yes	Yes

Reviewing REPAIR output

The potential output from the REPAIR utility consists of a modified page or pages in the specified DB2 table space or index, and a dump of the contents.

Error messages: At each LOCATE statement, the last data page and the new page being located are checked for a few common errors, and messages are issued.

Data checks: Although REPAIR enables you to manipulate both user and DB2 data by bypassing SQL, it does carry out some checking of data. For example, REPAIR is unable to write a page with the wrong page number, DB2 will abend with a 04E code and reason code C200B0. If the page is broken because the broken page bit is on or the incomplete page flag is set, REPAIR will issue the following message:

```
DSNU670I + DSNUCBRP - PAGE X'000004' IS A BROKEN PAGE
```

After running REPAIR

CHECK-pending status: You are responsible for violations of referential constraints caused by running REPAIR. These violations cause the target table space to be placed in the CHECK-pending status. See “Chapter 6. CHECK DATA” on page 53 for information about resetting this status.

Sample control statements

Example 1: Replacing damaged data and verifying replacement. Repair the specified page of table space DSN8S71E. Verify that, at the specified offset (50), the damaged data (A00) is found. Replace it with the desired data (D11). Initiate a dump beginning at offset 50, for 4 bytes, to verify the replacement.

```
//STEP1    EXEC  DSNUPROC,UID='IUIQU1UH',UTPROC='',SYSTEM='V7IA'
//SYSIN DD *
REPAIR OBJECT
  LOCATE TABLESPACE DSN8D71A.DSN8S71D PAGE X'02'
    VERIFY OFFSET 50 DATA X'A00'
    REPLACE OFFSET 50 DATA X'D11'
    DUMP OFFSET 50 LENGTH 4
```

Example 2: Removing a nonindexed row found by REORG. When reorganizing table space DSNDB04.TS1, you received the following message:

```
DSNU3401 DSNURBXA - ERROR LOADING INDEX, DUPLICATE KEY
                    INDEX = EMPINDEX
                    TABLE = EMP
                    RID OF INDEXED ROW  = X'00000201'
                    RID OF NONINDEXED ROW = X'00000503'
```

To resolve this error message, delete the nonindexed row and log the change. (The LOG keyword is not required; it is logged by default.)

```
REPAIR
  LOCATE TABLESPACE DSNDB04.TS1 RID (X'00000503')
    DELETE
```

Example 3: Report whether catalog and directory DBDs differ. Determine if the DBDs in the DB2 catalog and the DB2 directory are consistent for database DSN8D2AP.

```
REPAIR DBD TEST DATABASE DSN8D2AP
```

Example 4: Report differences between catalog and directory DBDs. After running the TEST option on database DSN8D2AP, and determining that the DBDs are inconsistent, determine the differences between the DBDs.

```
REPAIR DBD DIAGNOSE DATABASE DSN8D2AP OUTDDN SYSREC
```

Example 5: REPAIR table space with orphan row. After running DSN1CHKR on table space SYSDBASE, you received the following message:

```
DSN1812I ORPHAN ID = 20 ID ENTRY = 0190 FOUND IN
         PAGE = 000024
```

From a DSN1PRNT of page X'000024' and X'00002541', you identify that RID X'00002420' has a forward pointer of X'00002521'.

Repair the table space as follows:

1. Set the orphan's backward pointer to zeros.

```
REPAIR OBJECT LOG YES
 LOCATE TABLESPACE DSNDB06.SYSDBASE RID X'00002420
  VERIFY OFFSET X'0A' DATA X'00002422'
  REPLACE OFFSET X'0A' DATA X'00000000'
```

   Setting the pointer to zeros prevents the next step from updating link pointers while deleting, which can cause DB2 to abend if the orphan's pointers are incorrect.

2. Delete the orphan.

```
REPAIR OBJECT LOG YES
 LOCATE TABLESPACE DSNDB06.SYSDBASE RID X'00002420'
  VERIFY OFFSET X'06' DATA X'00002521'
  DELETE
```

Chapter 25. REPORT

The REPORT online utility provides information about table spaces. Use REPORT TABLESPACESET to find the names of all the table spaces and tables in a referential structure, including LOB table spaces. Use REPORT RECOVERY to find information necessary for recovering a table space, index, or a table space and all of its indexes. The REPORT utility also provides the LOB table spaces associated with a base table space.

For a diagram of REPORT syntax and a description of available options, see “Syntax and options of the control statement” on page 420. For detailed guidance on running this utility, see “Instructions for running REPORT” on page 423.

Output: The output from REPORT TABLESPACESET consists of the names of all table spaces in the table space set you specify. It also lists all tables in the table spaces and all tables dependent on those tables.

The output from REPORT RECOVERY consists of the recovery history from the SYSIBM.SYSCOPY catalog table, log ranges from the SYSIBM.SYSLGRNX directory table, and volume serial numbers where archive log data sets from the BSDS reside. In addition, REPORT RECOVERY output includes information about any indexes on the table space that are in the informational COPY-pending status, because this affects the recoverability of an index. For more information about this situation, see “Setting and clearing the informational COPY-pending status” on page 100.

In a data sharing environment, the output provides:

- The RBA when migrated to Version 7
- The high and low RBA values of the migrated member
- A list of any SYSLGRNX records from before data sharing was enabled that cannot be used to recover to any point in time after data sharing was enabled
- For SYSCOPY, the member that the image copy was deleted from

Authorization required: To execute this utility, the privilege set of the process must include one of the following:

- RECOVERDB privilege for the database
- DBADM or DBCTRL authority for the database
- SYSCTRL or SYSADM authority.

An ID with DBCTRL or DBADM authority over database DSNDB06 can run the REPORT utility on any table space in DSNDB01 (the directory) or DSNDB06 (the catalog), as can any ID with installation SYSOPR, SYSCTRL, or SYSADM authority.

Execution phases of REPORT: The REPORT utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup
REPORT	Information collection
UTILTERM	Cleanup

© Copyright IBM Corp. 1983, 2001

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

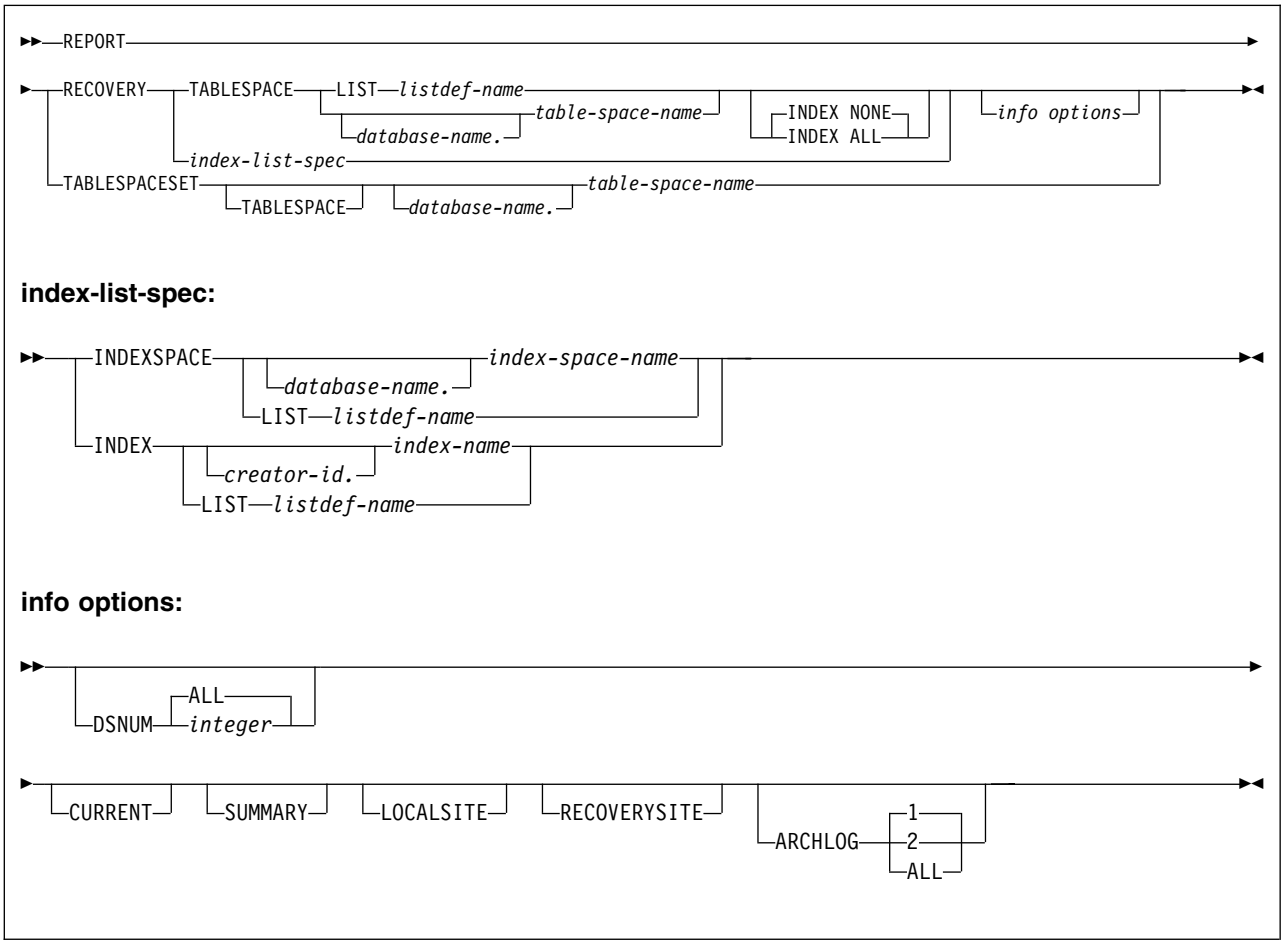

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

RECOVERY
Indicates that recovery information for the specified table space or index is to be reported.

TABLESPACESET
Indicates that the names of all table spaces in the table space set, as well as the names of all indexes over tables in the table space set are to be reported.

TABLESPACE *database-name.table-space-name*
For REPORT RECOVERY, specifies the table space (and, optionally, the database to which it belongs) being reported.

For REPORT TABLESPACESET, specifies a table space (and, optionally, the database to which it belongs) in the table space set.

database-name
Optionally specifies the database the table space belongs to.

table-space-name
Specifies the table space.

LIST *listdef-name*

Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each control statement of REPORT. The list must contain only table spaces. Do not specify LIST with the TABLESPACE ...*table-space-name* specification. The TABLESPACE keyword is required in order to validate the contents of the list. Utility list manager will invoke REPORT RECOVERY TABLESPACE once per item in the list.

For more information about LISTDEF specifications, see “Generating object lists using LISTDEF” on page 144.

INDEXSPACE *database-name.index-space-name*
Specifies the index space being reported.

database-name
Optionally specifies the database the index space belongs to.

index-space-name
Specifies the index space name for the index being reported.

LIST *listdef-name*

Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each control statement of REPORT. The list must contain only index spaces. Do not specify LIST with the INDEXSPACE...*index-space-name* specification. The INDEXSPACE keyword is required in order to validate the contents of the list. Utility list manager will invoke REPORT RECOVERY INDEXSPACE once for each item in the list.

For more information about LISTDEF specifications, see “Generating object lists using LISTDEF” on page 144.

INDEX *creator-id.index-name*
Specifies the index in the index space being reported.

creator-id
Optionally specifies the creator of the index.

index-name
Specifies the index name to be reported.

LIST *listdef-name*

Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each control statement of REPORT. The list must contain only index spaces. Do not specify LIST with the INDEX...*index-name* specification. The INDEX keyword is required in

order to validate the contents of the list. Utility list manager will invoke REPORT RECOVERY INDEX once for each item in the list.

For more information about LISTDEF specifications, see "Generating object lists using LISTDEF" on page 144.

The following keywords are optional:

INDEX NONE
Does not report recovery information for index spaces associated with the specified table space.

INDEX ALL
Reports recovery information for index spaces associated with the specified table space.

DSNUM
For a table space, identifies a partition or data set, for which information is to be reported; or it reports information for the entire table space. For an index space, identifies a partition for which information is to be reported; or it reports information for the entire index space.

ALL Reports information for the entire table space or index space. The **default** is **ALL**.

integer
Is the number of a partition or data set for which information is to be reported. The maximum is 254.

For a partitioned table space or index space, the integer is its partition number.

For a nonpartitioned table space, find the integer at the end of the data set name as cataloged in the VSAM catalog. The data set name has this format:

catname.DSNDB*x*.*dbname*.*tsname*.*y*0001.A*nnn*

where:

catname	The VSAM catalog name or alias
x	C or D
dbname	The database name
tsname	The table space name
y	I or J
nnn	The data set integer.

CURRENT
Specifies that only the SYSCOPY entries after the last recoverable point of the table space are reported. The last recoverable point is the last full image copy, LOAD REPLACE LOG YES or REORG LOG YES. If you specify DSNUM ALL, then the last recoverable point is a full image copy that was taken for the entire table space or index space. However, if you specify the CURRENT option, but the last recoverable point does not exist on the active log, DB2 prompts you to mount archive tapes until this point is found.

CURRENT also reports only the SYSLGRNX rows and archive log volumes after the last incremental image copy entry. If no incremental image copies were created, only the SYSLGRNX rows and archive log volumes created after the last recoverable point are reported.

If you do not specify CURRENT or if no last recoverable point exists, all SYSCOPY and SYSLGRNX entries for that table space or index space are

reported, including those on archive logs. If you do not specify CURRENT, the entries after the last recoverable point are marked with an asterisk (*) in the report.

SUMMARY
Specifies that only a summary of volume serial numbers is to be reported. It reports the following volume serial numbers:
- Where the archive log data sets from the BSDS reside
- Where the image copy data sets from SYSCOPY reside.

If you do not specify SUMMARY, recovery information is reported in addition to the summary of volume serial numbers.

LOCALSITE
Specifies a report of all SYSCOPY records copied from a local site system.

RECOVERYSITE
Specifies a report of all SYSCOPY records copied for the recovery site system.

ARCHLOG
Specifies which archive log data sets are to be reported.

1 Reports archive log data set 1 only.

The **default** is **1**.

2 Reports archive log data set 2 only.

ALL
Optionally reports both archive log data sets 1 and 2.

Instructions for running REPORT

To run REPORT, you must:

1. Prepare the necessary data sets, as described in "Data sets used by REPORT".
2. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for REPORT, see "Sample control statements" on page 429.)
3. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 424. (For a complete description of the syntax and options for REPORT, see "Syntax and options of the control statement" on page 420.)
4. Check the compatibility table in "Concurrency and compatibility" on page 425 if you want to run other jobs concurrently on the same target objects.
5. Plan for restart if the REPORT job doesn't complete, as described in "Terminating or restarting REPORT" on page 425.
6. Run REPORT.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Data sets used by REPORT

Table 82 on page 424 describes the data sets used by REPORT. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 82. Data sets used by REPORT

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table space
Object to be reported. It is named in the REPORT control statement and is accessed through the DB2 catalog.

Creating the control statement

See "Syntax and options of the control statement" on page 420 for REPORT syntax and option descriptions. See "Sample control statements" on page 429 for examples of REPORT usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:
"Reporting recovery information"
"Running REPORT on the catalog and directory" on page 425

Reporting recovery information

You can use the REPORT utility in planning for recovery. REPORT provides information necessary for recovering a table space. You can request report information for LOCALSITE or RECOVERYSITE, or both. REPORT RECOVERY displays:

- Recovery information from the SYSIBM.SYSCOPY catalog table, including QUIESCE, COPY, LOAD, REORG, RECOVER TOCOPY, and RECOVER TORBA (or TOLOGPOINT) history. It also indicates device type and whether this is the primary or backup copy for LOCALSITE or RECOVERYSITE.
- Log ranges of the table space from the SYSIBM.SYSLGRNX directory.
- Archive log data sets ARCHLOG1, ARCHLOG2, or both from the bootstrap data set.

You can use REPORT TABLESPACESET to find the names of all members of a table space set.

You can also use REPORT to obtain recovery information about the catalog and directory. When doing so, use the CURRENT option to avoid unnecessary mounting of archive tapes.

REPORT denotes any non-COPY entries it finds in the SYSCOPY catalog table with asterisks. For example, an entry added by the QUIESCE utility is marked with asterisks in the REPORT output.

The following statement reports the names of all table spaces in the table space set containing table space DSN8S71E:

```
REPORT TABLESPACESET TABLESPACE DSN8D71A.DSN8S71E
```

The following statement reports recovery information for table space DSN8S71D for the local subsystem only:

```
REPORT RECOVERY TABLESPACE DSN8D71A.DSN8S71D LOCALSITE
```

For image copies of partitioned table spaces taken with the DSNUM ALL option, we recommend that you run REPORT RECOVERY DSNUM ALL. If you run REPORT RECOVERY DSNUM ALL CURRENT, DB2 reports additional historical information dating back to the last full image copy taken for the entire table space.

The REPORT RECOVERY utility output indicates if any image copies are unusable; image copies taken prior to REORG or LOAD events to reset REORG-pending status are marked as unusable. REPORT output indicates which image copies will reset the REORG-pending status by displaying the ICTYPE field as <R>, as shown in Figure 23.

```
DSNU582I = DSNUPPCP - REPORT RECOVERY TABLESPACE LDB1.TS1 SYSCOPY ROWS
| TIMESTAMP = 2000-11-15-07.47.16.860892,  IC TYPE = <R>, SHR LVL = ,  DSNUM    = 0000,  START LRSN =B4F2FEE2A33A
  DEV TYPE  =         ,                    IC BACK =   ,  STYPE   = A, FILE SEQ = 0000,  PIT LRSN  = 000000000000
  LOW DSNUM = 0000,                        HIGH DSNUM = 0000
  DSNAME    = LDB1.TS1                                    ,   MEMBER NAME =
  LOWDSNUM = 0001 HIGHDSNUM = 0005
```

Figure 23. Example of REPORT RECOVERY unusable copy indication

Running REPORT on the catalog and directory

REPORT RECOVERY shows the image copies for those table spaces that are not included in SYSIBM.SYSCOPY:

- DSNDB01.SYSUTILX
- DSNDB01.DBD01
- DSNDB06.SYSCOPY

When you execute REPORT RECOVERY on DSNDB01.DBD01, DSNDB01.SYSUTIL, or DSNDB06.SYSCOPY, specify the CURRENT option to avoid unnecessarily mounting archive tapes. If you do not specify CURRENT, DB2 searches for and reports all SYSCOPY records in the log, including those on archive tapes. However, if the CURRENT option is specified and the last recoverable point does not exist on the active log, DB2 prompts you to mount archive tapes until this point is found.

REPORT TABLESPACESET can be used on the DB2 catalog and directory table spaces.

Terminating or restarting REPORT

You can restart a REPORT utility job, but it starts from the beginning again.

You can terminate REPORT with the TERM UTILITY command.

For guidance in restarting online utilities, see “Restarting an online utility” on page 46.

Concurrency and compatibility

REPORT does not set a utility restrictive state on the target table space or partition.

REPORT can run concurrently on the same target object with any utility or SQL operation.

Reviewing REPORT output

REPORT TABLESPACESET output: The output from REPORT TABLESPACESET consists of the names of all table spaces in the table space set you specify. It also specifies all tables in the table spaces, and specifies all tables dependent on those tables, including LOB table spaces.

For example, REPORT TABLESPACESET TABLESPACE LDB1.TS1 generates the following output:

```
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = RECTS004.CFD1
DSNU050I   DSNUGUTC - REPORT TABLESPACESET TABLESPACE LDB1.TS1
DSNU587I - DSNUPSET - REPORT TABLESPACE SET WITH TABLESPACE LDB1.TS1
                  TABLESPACE            TABLE                           DEPENDENT TABLE
                  LDB1.TS1              SYSADM.T1
                                        SYSADM.T2
DSNU580I   DSNUPORT - REPORT UTILITY COMPLETE - ELAPSED TIME=00:00:00
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

Figure 24. Example of REPORT TABLESPACESET

REPORT RECOVERY output: REPORT RECOVERY displays all the information about the image copy data sets and archive log data set that might be required during the recover.

If the DSVOLSER column of SYSIBM.SYSCOPY is blank, REPORT RECOVERY does not display volume serial numbers for image copy data sets.

The report contains 3 sections, which include the following types of information:

- Recovery history from the SYSIBM.SYSCOPY catalog table.

  For a description of the fields in the SYSCOPY rows, see the table describing SYSIBM.SYSCOPY in Appendix D of *DB2 SQL Reference*.
- Log ranges from SYSIBM.SYSLGRNX.
- Volume serial numbers where archive log data sets from the BSDS reside.

If there is no data to report for one or more of these topics, the corresponding sections of the report contain this message:

```
DSNU588I - NO DATA TO BE REPORTED
```

Figure 25 on page 427 is a sample of REPORT RECOVERY output in a data sharing environment.

```
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = RECTS004.CFD1
DSNU050I   DSNUGUTC - REPORT RECOVERY TABLESPACE LDB1.TS1
DSNU581I - DSNUPREC - REPORT RECOVERY TABLESPACE LDB1.TS1
DSNU593I - DSNUPREC -  REPORT RECOVERY ENVIRONMENT RECORD:
'             MINIMUM   RBA: 000000000000
'             MAXIMUM   RBA: 000001DD9AE8
'             MIGRATING RBA: 000001DD9AE8
DSNU582I - DSNUPPCP - REPORT RECOVERY TABLESPACE LDB1.TS1 SYSCOPY ROWS
TIMESTAMP = 1998-09-10-08.21.25.912161,  IC TYPE = *Y*,  SHR LVL =  ,  DSNUM    = 0000,  START LRSN =AAE19AEEE0D4
DEV TYPE  =         ,                    IC BACK =   ,   STYPE   =  ,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = LDB1.TS1                                    ,   MEMBER NAME = DB1G

TIMESTAMP = 1998-09-10-08.21.46.464782,  IC TYPE =  F ,  SHR LVL = R,  DSNUM    = 0000,  START LRSN =AAE19B0B107B
DEV TYPE  = 3390    ,                    IC BACK =   ,   STYPE   =  ,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = CPYLPF1                                     ,   MEMBER NAME = DB1G

TIMESTAMP = 1998-09-10-08.22.07.674391,  IC TYPE = *Q*,  SHR LVL =  ,  DSNUM    = 0000,  START LRSN =AAE19B1EBCAF
DEV TYPE  =         ,                    IC BACK =   ,   STYPE   = W,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = LDB1.TS1                                    ,   MEMBER NAME = DB1G

TIMESTAMP = 1998-09-10-08.22.15.786373,  IC TYPE =  I ,  SHR LVL = R,  DSNUM    = 0000,  START LRSN =AAE19B24CC8A
DEV TYPE  = 3390    ,                    IC BACK =   ,   STYPE   =  ,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = CPYLPF2                                     ,   MEMBER NAME = DB1G

TIMESTAMP = 1998-09-10-08.22.22.810734,  IC TYPE = *Q*,  SHR LVL =  ,  DSNUM    = 0000,  START LRSN =AAE19B3C986B
DEV TYPE  =         ,                    IC BACK =   ,   STYPE   = W,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = LDB1.TS1                                    ,   MEMBER NAME = DB1G

TIMESTAMP = 1998-09-10-08.22.35.623910,  IC TYPE =  I ,  SHR LVL = R,  DSNUM    = 0000,  START LRSN =AAE19B4342EC
DEV TYPE  = 3390    ,                    IC BACK =   ,   STYPE   =  ,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = CPYLPF3                                     ,   MEMBER NAME = DB1G
DSNU586I - DSNUPSUM - REPORT RECOVERY TABLESPACE LDB1.TS1 SUMMARY
DSNU588I - DSNUPSUM - NO DATA TO BE REPORTED

DSNU592I - DSNUPREC - REPORT RECOVERY INFORMATION FOR DATA SHARING MEMBER : DB1G
DSNU583I - DSNUPPLR - SYSLGRNX ROWS FROM REPORT RECOVERY FOR TABLESPACE LDB1.TS1
 UCDATE     UCTIME      START RBA      STOP  RBA     START LRSN     STOP  LRSN    PARTITION   MEMBER ID
 091098    16375569   000002017A57   000002019033   AAE19AD84EB1   AAE19AD8BA67     0000         0001
 091098    16380534   00000201C61C   000002022F16   AAE19AE1822F   AAE19AEE6DE6     0000         0001
 091098    16382567   AAE19AF4653C   AAE19AF4653C   AAE19AF4653C   AAE19AF4653C     0000         0001
 091098    16390230   0000020297AB   00000202FA14   AAE19B17C9D3   AAE19B1E7E32     0000         0001
 091098    16392787   000002034448   0000020663CE   AAE19B3037E6   AAE19B37FE01     0000         0001


DSNU586I - DSNUPSUM - REPORT RECOVERY TABLESPACE LDB1.TS1 SUMMARY
DSNU588I - DSNUPSUM - NO DATA TO BE REPORTED

DSNU592I - DSNUPREC - REPORT RECOVERY INFORMATION FOR DATA SHARING MEMBER : DB2G



DSNU586I - DSNUPSUM - REPORT RECOVERY TABLESPACE LDB1.TS1 SUMMARY
DSNU588I - DSNUPSUM - NO DATA TO BE REPORTED

DSNU592I - DSNUPREC - REPORT RECOVERY INFORMATION FOR DATA SHARING MEMBER : DB3G
```

Figure 25. Example of REPORT RECOVERY (Part 1 of 2)

REPORT

```
DSNU586I - DSNUPSUM - REPORT RECOVERY TABLESPACE LDB1.TS1 SUMMARY
DSNU588I - DSNUPSUM - NO DATA TO BE REPORTED
DSNU589I - DSNUPREC - REPORT RECOVERY TABLESPACE LDB1.TS1 COMPLETE
DSNU580I   DSNUPORT - REPORT UTILITY COMPLETE - ELAPSED TIME=00:00:00
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

Figure 25. Example of REPORT RECOVERY (Part 2 of 2)

Figure 26 is a sample of REPORT RECOVERY TABLESPACE ARCHLOG output. Under message DSNU584I, the archive log entries after the last recoverable point are marked with an asterisk (*). If you code the CURRENT option, the output from message DSNU584I would include only the archive logs after the last recoverable point and the asterisk (*) would not be included in the report.

```
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = REPORT
DSNU050I   DSNUGUTC - REPORT RECOVERY TABLESPACE MYDBASE.MYSPACE ARCHLOG ALL
DSNU581I - DSNUPREC - REPORT RECOVERY TABLESPACE MYDBASE.MYSPACE
DSNU593I - DSNUPREC -  REPORT RECOVERY ENVIRONMENT RECORD:
'             MINIMUM   RBA: 000000000000
'             MAXIMUM   RBA: FFFFFFFFFFFF
'             MIGRATING RBA: 000000000000
DSNU582I - DSNUPPCP - REPORT RECOVERY TABLESPACE MYDBASE.MYSPACE SYSCOPY ROWS
TIMESTAMP = 1998-12-11-09.38.21.734394,  IC TYPE = *Y*,  SHR LVL =  ,  DSNUM    = 0000,  START LRSN =00000264B9A
DEV TYPE  =         ,                    IC BACK =   ,   STYPE   =  ,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = MYCOPYA                                    ,   MEMBER NAME =
DS VOLSER = SCR03 ,

TIMESTAMP = 1998-12-11-09.38.48.913881,  IC TYPE =  F ,  SHR LVL = R,  DSNUM    = 0000,  START LRSN =00000266354
DEV TYPE  = 3390    ,                    IC BACK =   ,   STYPE   =  ,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = MYCOPYIA                                   ,   MEMBER NAME = DB1G
DS VOLSER = SCR03 ,

TIMESTAMP = 1998-12-11-09.39.09.542154,  IC TYPE = *Q*,  SHR LVL =  ,  DSNUM    = 0000,  START LRSN =000002674A2
DEV TYPE  =         ,                    IC BACK =   ,   STYPE   = W,  FILE SEQ = 0000,  PIT LRSN  = 000000000000
LOW DSNUM = 0000,                        HIGH DSNUM = 0000
DSNAME    = MYCOPYIB                                   ,   MEMBER NAME = DB1G
DS VOLSER = SCR03 ,

DSNU586I - DSNUPSUM - REPORT RECOVERY TABLESPACE MYDBASE.MYSPACE SUMMARY
               IC COPY VOLSER(S)             SCR03

DSNU583I - DSNUPPLR - SYSLGRNX ROWS FROM REPORT RECOVERY FOR TABLESPACE MYDBASE.MYSPACE

      START RBA      STOP  RBA     START LRSN     STOP  LRSN    PARTITION   MEMBER ID
    00000262E3A6   00000262E86F   00000262E3A6   00000262E86F     0000         0000
    000002632CC8   0000026394D4   000002632CC8   0000026394D4     0000         0000
    000002652EEF   0000026535A1   000002652EEF   0000026535A1     0000         0000
    000002671E1C   0000026724B9   000002671E1C   0000026724B9     0000         0000
    0000026792AA   000002679B08   0000026792AA   000002679B08     0000         0000

DSNU584I - DSNUPPBS - REPORT RECOVERY TABLESPACE MYDBASE.MYSPACE ARCHLOG1 BSDS VOLUMES

      START RBA      END RBA       UNIT  VOLSER   DATA SET NAME
    000002328000   00000263DFFF   SYSDA  SCR03  DSNC420.ARCHLOG1.A0000003
    00000263E000   000002643FFF   SYSDA  SCR03  DSNC420.ARCHLOG1.A0000004
    000002644000   000002648FFF   SYSDA  SCR03  DSNC420.ARCHLOG1.A0000005  *
    000002649000   00000264CFFF   SYSDA  SCR03  DSNC420.ARCHLOG1.A0000006  *
    00000264D000   000002653FFF   SYSDA  SCR03  DSNC420.ARCHLOG1.A0000007  *
    000002654000   000002658FFF   SYSDA  SCR03  DSNC420.ARCHLOG1.A0000008  *
    000002659000   00000265EFFF   SYSDA  SCR03  DSNC420.ARCHLOG1.A0000009  *
    00000265F000   000002666FFF   SYSDA  SCR03  DSNC420.ARCHLOG1.A0000010  *
    000002667000   00000266BFFF   SYSDA  SCR04  DSNC420.ARCHLOG1.A0000011  *
    00000266C000   00000267BFFF   SYSDA  SCR04  DSNC420.ARCHLOG1.A0000012  *
```

Figure 26. Example of REPORT RECOVERY TABLESPACE ARCHLOG (Part 1 of 2)

```
DSNU5834I - DSNUPPBS - REPORT RECOVERY TABLESPACE MYDBASE.MYSPACE ARCHLOG2 BSDS VO

      START RBA      END RBA       UNIT  VOLSER   DATA SET NAME
    000002328000   00000263DFFF   SYSDA  SCR03  DSNC420.ARCHLOG2.A0000003
    00000263E000   000002643FFF   SYSDA  SCR03  DSNC420.ARCHLOG2.A0000004
    000002644000   000002648FFF   SYSDA  SCR03  DSNC420.ARCHLOG2.A0000005  *
    000002649000   00000264CFFF   SYSDA  SCR03  DSNC420.ARCHLOG2.A0000006  *
    00000264D000   000002653FFF   SYSDA  SCR03  DSNC420.ARCHLOG2.A0000007  *
    000002654000   000002658FFF   SYSDA  SCR03  DSNC420.ARCHLOG2.A0000008  *
    000002659000   00000265EFFF   SYSDA  SCR03  DSNC420.ARCHLOG2.A0000009  *
    00000265F000   000002666FFF   SYSDA  SCR03  DSNC420.ARCHLOG2.A0000010  *
    000002667000   00000266BFFF   SYSDA  SCR04  DSNC420.ARCHLOG2.A0000011  *
    00000266C000   00000267BFFF   SYSDA  SCR04  DSNC420.ARCHLOG2.A0000012  *

SNU586I - DSNUPSUM - REPORT RECOVERY TABLESPACE MYDBASE.MYSPACE SUMMARY
              ARCHLOG1 BSDS VOLSERS(S)      SCR03 *
                                            SCR04 *
              ARCHLOG2 BSDS VOLSERS(S)      SCR03 *
                                            SCR04 *

DSNU589I - DSNUPREC - REPORT RECOVERY TABLESPACE MYDBASE.MYSPACE COMPLETE
DSNU580I   DSNUPORT - REPORT UTILITY COMPLETE - ELAPSED TIME=00:00:00
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

Figure 26. Example of REPORT RECOVERY TABLESPACE ARCHLOG (Part 2 of 2)

Sample control statements

Example 1: Sample JCL for REPORT RECOVERY.

```
//STEP1    EXEC DSNUPROC,UID='IUKUU206.REPORT2',
//            UTPROC='',SYSTEM='DSN'
//SYSIN    DD *
REPORT RECOVERY
       TABLESPACE DSN8D71A.DSN8S71E
//*
```

Example 2: Sample control statement for REPORT TABLESPACESET.

```
REPORT TABLESPACESET
       TABLESPACE UTQPD22A.UTQPS22E
```

Example 3: REPORT referentially related table spaces. The following statement reports the names of all table spaces in the table space set containing table space DSN8D71A.DSN8S71E.

```
REPORT TABLESPACESET TABLESPACE DSN8D71A.DSN8S71E
```

Example 4: REPORT RECOVERY information for a table space. This statement reports recovery information for table space DSN8D71A.DSN8S71D.

```
REPORT RECOVERY TABLESPACE DSN8D71A.DSN8S71D DSNUM ALL
```

REPORT

Chapter 26. RUNSTATS

The RUNSTATS online utility gathers summary information about the characteristics of data in table spaces, indexes, and partitions. DB2 records this information in the DB2 catalog and uses it to select access paths to data during the bind process. It is available to the database administrator for evaluating database design and to aid in determining when table spaces or indexes must be reorganized.

There are two formats for the RUNSTATS utility: RUNSTATS TABLESPACE and RUNSTATS INDEX. RUNSTATS TABLESPACE gathers statistics on a table space and, optionally, on indexes or columns; RUNSTATS INDEX gathers statistics only on indexes.

Use the STATISTICS keyword with LOAD, REBUILD INDEX, and REORG jobs to eliminate the need to execute RUNSTATS for updating catalog statistics. If you restart a LOAD or REBUILD INDEX job that uses the STATISTICS keyword, inline statistics collection does not occur. To update catalog statistics, run the RUNSTATS utility after the restarted utility job completes. For information about restarting a REORG job which uses the STATISTICS keyword, see page 382.

Use the HISTORY keyword to collect the statistics history. When RUNSTATS stores new information in the DB2 catalog tables, it will also store that information in the V7 catalog history tables. The catalog will hold the current statistics while the catalog history tables will hold historical, including current, statistics.

You can specify that a LOB table space is to have space statistics collected so you can determine when the LOB table space should be reorganized. No statistics on the LOB table space affect access path selection.

DB2 invalidates statements in the dynamic statement cache when you run RUNSTATS against objects to which those statements refer. In a data sharing environment, the relevant statements are also invalidated in the cache of other members in the group. DB2 invalidates the cached statements to ensure that the next invocations of those statements are fully prepared and pick up the latest access path changes.

For a diagram of RUNSTATS syntax and a description of available options, see “Syntax and options of the control statement” on page 432. For detailed guidance on running this utility, see “Instructions for running RUNSTATS” on page 441.

Output: RUNSTATS updates the DB2 catalog with table space or index space statistics or prints a report. The information updated by RUNSTATS is used by DB2 to select access paths to the data. You can query the catalog tables to obtain the updated statistics. See “Reviewing RUNSTATS output” on page 445 for a list of all the catalog tables and columns updated by RUNSTATS.

Additional information:
- STATS privilege for the database
- DBADM, DBCTRL, or DBMAINT authority for the database
- SYSCTRL or SYSADM authority.

An ID with installation SYSOPR authority can also execute RUNSTATS, but only on a table space in the DSNDB06 database.

© Copyright IBM Corp. 1983, 2001

To use REPORT YES, you must have the SELECT privilege on the tables reported. Values are not reported from the tables the user is not authorized to see.

Execution phases of RUNSTATS: The RUNSTATS utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup
RUNSTATS	Scanning table space or index and updating catalog
UTILTERM	Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

RUNSTATS TABLESPACE syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

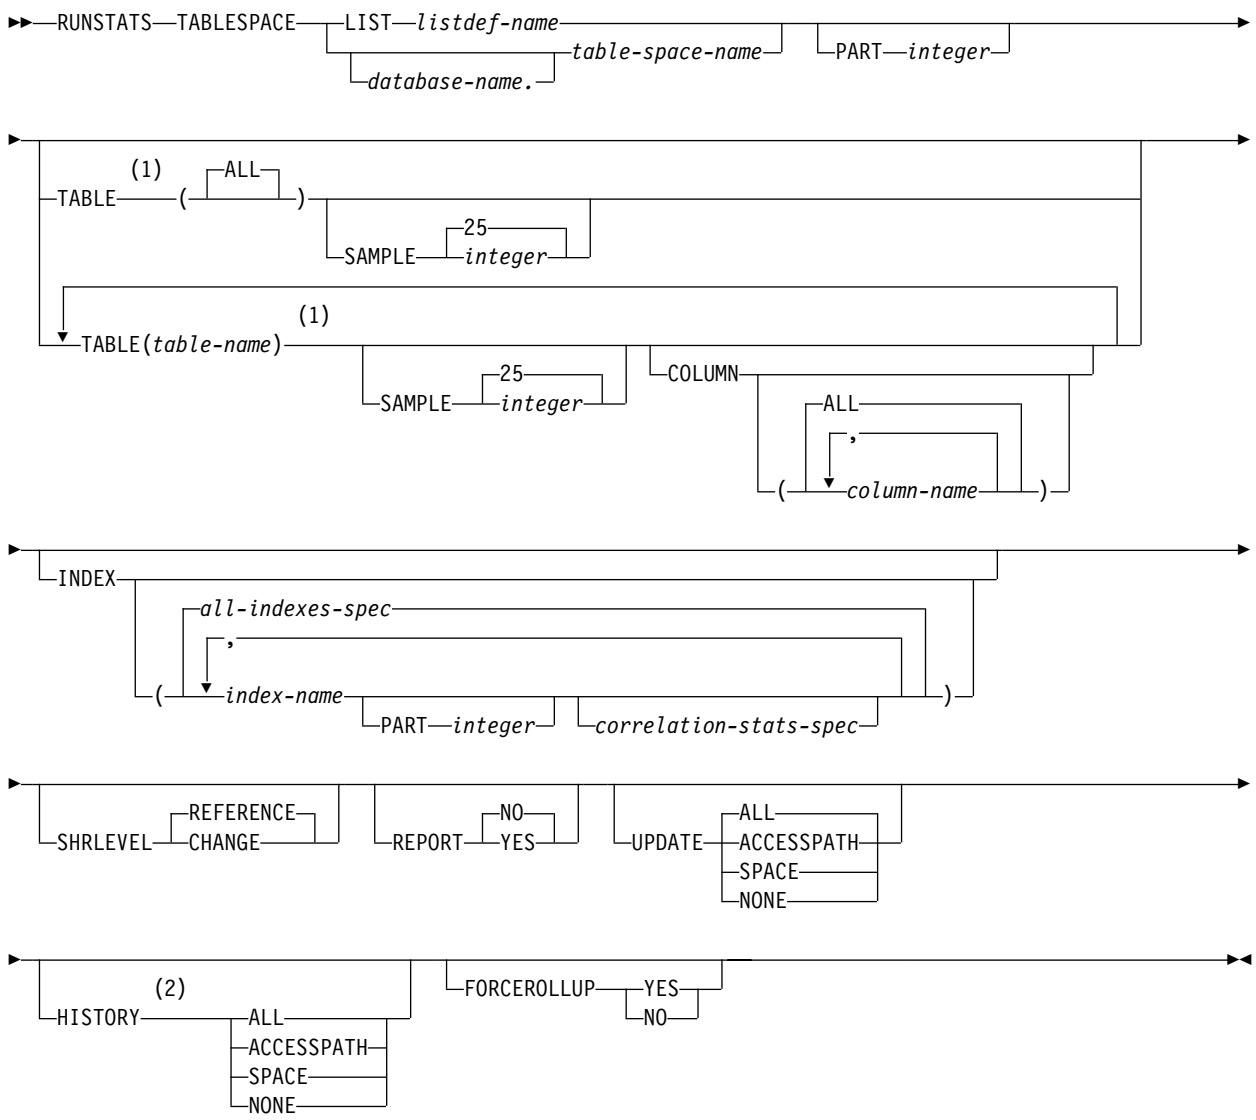

Notes:

1 The TABLE keyword is not valid for a LOB table space.

2 What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

all-indexes-spec:

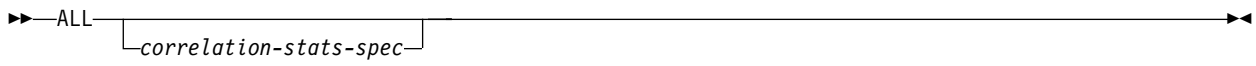

correlation-stats-spec:

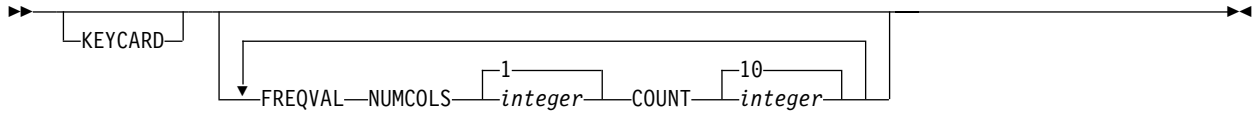

RUNSTATS TABLESPACE option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

TABLESPACE *database-name.table-space-name*
Specifies the table space (and, optionally, the database to which it belongs) for which table space and table information is to be gathered. It must not be a table space in DSNDB01 or DSNDB07.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each control statement of RUNSTATS. The list must contain only table spaces for RUNSTATS TABLESPACE. Do not specify LIST with keywords from the **TABLE**...(*table-name*) specification. Instead, specify LIST with TABLE (ALL).The same is true of INDEX(ALL). You cannot specify index names with a list. Use INDEX(ALL).

Also, PART can not be specified. Use LISTDEF PARTLEVEL instead. The TABLESPACE keyword is required in order to validate the contents of the list. Utility list manager will invoke RUNSTATS TABLESPACE once for each item in the list.

For more information about LISTDEF specifications, see “Generating object lists using LISTDEF” on page 144.

database-name
Is the name of the database to which the table space belongs.

The **default** is **DSNDB04**.

table-space-name
Is the name of the table space about which information is gathered.

If the table space specified by the TABLESPACE keyword is a LOB table space, these are the only keywords allowed: SHRLEVEL REFERENCE or CHANGE, REPORT YES or NO, and UPDATE ALL or NONE.

PART *integer*
Identifies a table space partition for which statistics are to be collected.

integer is the number of the partition and must be in the range from 1 to the number of partitions defined for the table space. The maximum is 254.

PART cannot be specified with LIST.

TABLE
Specifies the table for which column information is to be gathered. All tables must belong to the table space specified in the TABLESPACE option.

(ALL) Specifies that information is to be gathered for all columns of all tables in the table space.

The **default** is **ALL**.

The TABLE option is not valid for a LOB table space.

SAMPLE *integer*
Indicates the percentage of rows to sample when collecting non-indexed column statistics. Any value from 1 through 100 can be specified. The **default** is **25**.

The SAMPLE option is not allowed for LOB table spaces.

(*table-name***)**
Specifies the tables for which column information is to be gathered. The parentheses are required. If you omit the qualifier, the user identifier for the utility job is used.

If you specify more than one table, you must repeat the TABLE option.

COLUMN
Specifies columns for which column information is to be gathered.

You can only specify this option if you specify a particular tables for which statistics are to be gathered (TABLE (*table-name*)). If you specify particular tables and do not specify the COLUMN option, the default, **COLUMN**(**ALL**), is used. If you do not specify a particular table when using the TABLE option, you cannot specify the COLUMN option; however, COLUMN(ALL) is assumed.

(ALL)
Specifies that statistics are to be gathered for all columns in the table.

The **COLUMN (ALL)** option is not allowed for LOB table spaces.

(*column-name***, ...)**
Specifies the columns for which statistics are to be gathered. The parentheses are required.

You can specify a list of column names; the maximum is 10. If you specify more than one column, separate each name with a comma.

INDEX
Specifies indexes for which information is to be gathered. Column information is gathered for the first column of the index, and might be gathered for additional index columns depending on the options you specify. All the indexes must be associated with the *same* table space, which must be the table space specified in the TABLESPACE option.

INDEX is valid for gathering statistics on an index on the auxiliary table.

(ALL) Specifies that the column information is to be gathered for all indexes defined on tables contained in the table space. The parentheses are required.

The **default** is **ALL**.

(*index-name***, ...)**
Specifies the indexes for which information is to be gathered. The parentheses are required.

You can specify a list of index names. If you specify more than one index, separate each name with a comma.

PART *integer*
Identifies an index partition for which statistics are to be collected.

integer is the number of the partition.

SHRLEVEL
Tells whether other programs that access the table space while RUNSTATS is running must use read-only access, or can change the table space.

REFERENCE
Allows only read-only access by other programs.

The **default** is **REFERENCE**.

CHANGE
Allows other programs to change the table space or index. With SHRLEVEL CHANGE, uncommitted data can be collected into statistical summaries.

REPORT
Determines if a set of messages is generated to report the collected statistics.

NO
Indicates that the set of messages is not output to SYSPRINT.

The **default** is **NO**.

YES
Indicates that the set of messages is output to SYSPRINT. The messages generated are dependent on the combination of keywords (such as TABLESPACE, INDEX, TABLE, and COLUMN) specified with the RUNSTATS utility. However, these messages are *not* dependent on the specification of the UPDATE option. REPORT YES always generates a report of SPACE and ACCESSPATH statistics.

UPDATE
Tells whether the collected statistics are inserted into the catalog tables. UPDATE also allows you to select statistics used for access path selection or statistics used by database administrators. What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

ALL Indicates that all collected statistics will be updated in the catalog.

The **default** is **ALL**.

ACCESSPATH
Indicates that only the catalog table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only the catalog table columns that provide statistics to help the database administrator assess the status of a particular table space or index are updated.

NONE Indicates that no catalog tables are updated with the collected statistics. This option is only valid when REPORT YES is specified.

HISTORY
Records all catalog table inserts or updates to the catalog history tables. What you specify under the UPDATE option affects what options you can specify under the HISTORY option. The table below lists the HISTORY options allowed for each UPDATE option specified.

The default will be supplied by the value specified in STATISTICS HISTORY on panel DSNTIPO.

ALL Indicates that all collected statistics will be updated in the catalog history tables.

ACCESSPATH
Indicates that only the catalog history table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only SPACE-related catalog statistics are updated in catalog history tables.

NONE Indicates that no catalog history tables are updated with the collected statistics.

Table 83. Allowable HISTORY options based on selected UPDATE option

UPDATE option	HISTORY options
ALL	ALL, ACCESSPATH, SPACE, NONE
ACCESSPATH	ACCESSPATH, NONE
SPACE	SPACE, NONE
NONE	NONE

FORCEROLLUP
Determines if aggregation or rollup of statistics will take place when RUNSTATS is executed in spite of some parts being empty. This will enable the optimizer to select the best access path.

The following are the available options for the FORCEROLLUP Keyword.

YES Indicates that forced aggregation or rollup processing is to be done, even though some parts may not contain data.

NO Indicates that aggregation or rollup will be done only if data is available for all parts.

If data is not available for all parts, DSNU623I message will be issued, provided the installation value for STATISTICS ROLLUP on panel DSNTIPO is set to NO.

KEYCARD
Collects all of the distinct values in all of the 1 to *n* key column combinations for the specified indexes. *n* is the number of columns in the index.

FREQVAL
Controls the collection of frequent value statistics. If you specify FREQVAL, it must be followed by two additional keywords:

NUMCOLS
Indicates the number of key columns to concatenate together when collecting frequent values from the specified index. Specifying '3' means to collect frequent values on the concatenation of the first three key columns. The **default** is **1**, which means collect frequent values on the first key column of the index.

COUNT
Indicates the number of frequent values to be collected. Specifying '15' means collect 15 frequent values from the specified key columns. The **default** is **10**.

RUNSTATS INDEX syntax diagram

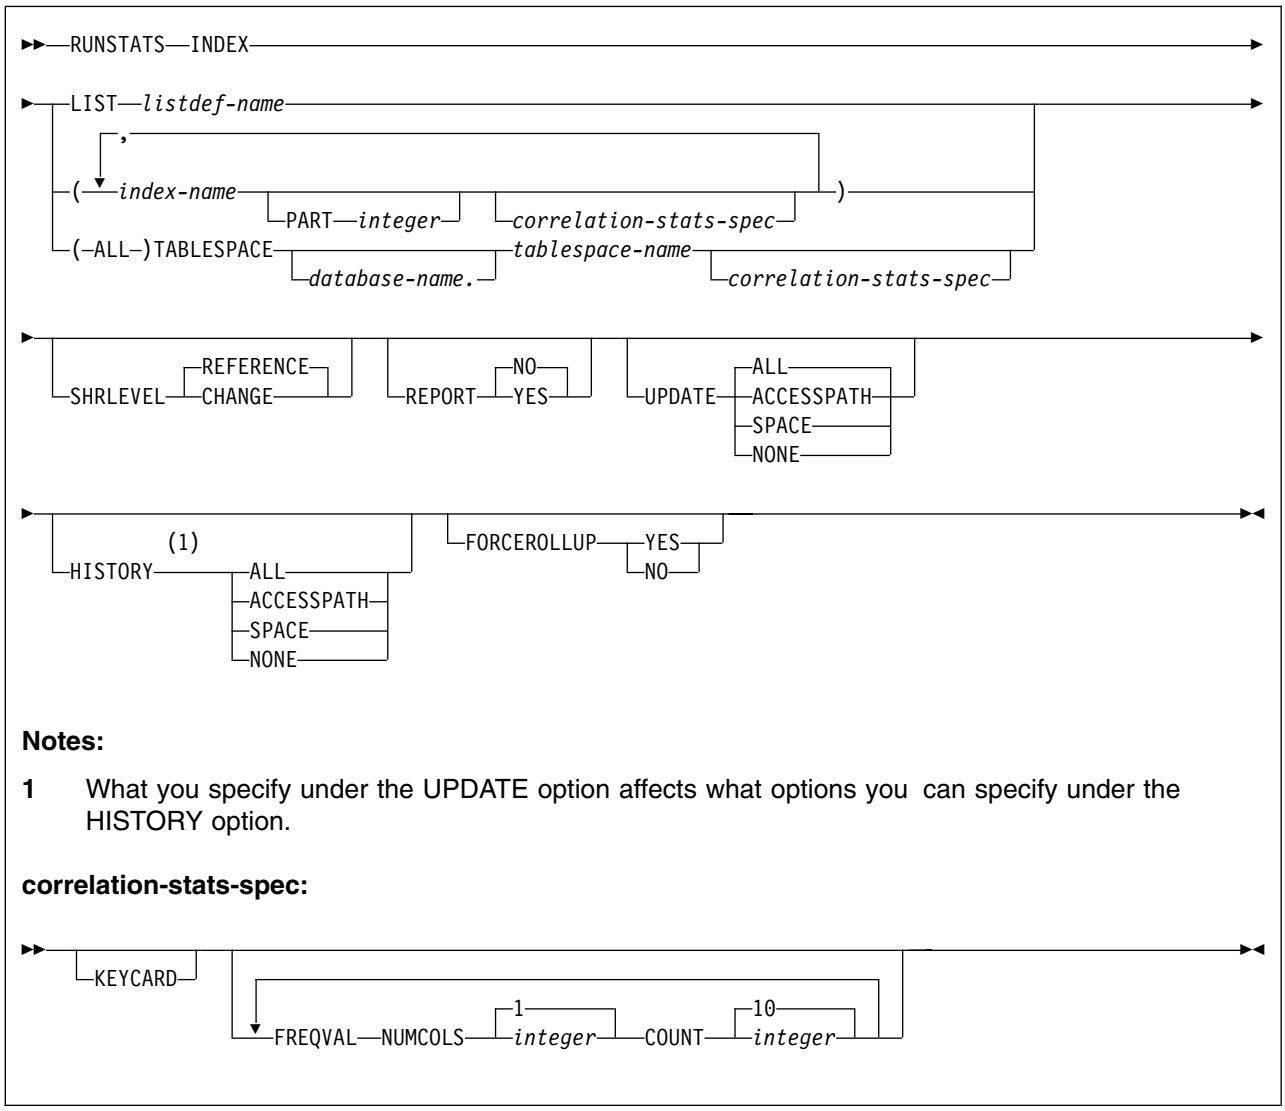

Notes:

1 What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

correlation-stats-spec:

RUNSTATS INDEX option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

INDEX
Specifies indexes for which information is to be gathered. Column information is gathered for the first column of the index. All the indexes must be associated with the *same* table space.

LIST *listdef-name*
Specifies the name of a previously defined LISTDEF list name. The utility allows one LIST keyword for each control statement of RUNSTATS. The list must contain only index spaces for RUNSTATS INDEX. Do not specify LIST with keywords from the **INDEX**...(*index-name*) specification.

The utility list manager will group indexes by their related table space and execute runstats index once per table space. The INDEX keyword is required in order to validate the contents of the LIST.

For more information about LISTDEF specifications, see “Generating object lists using LISTDEF” on page 144.

(*index-name***, ...)**
Specifies the indexes for which information is to be gathered. The parentheses are required.

You can specify a list of index names. If you specify more than one index, separate each name with a comma.

PART *integer*
Identifies the index partition for which statistics are to be collected.

integer is the number of the partition.

(ALL)
Specifies that information is to be gathered for all indexes defined on all tables in the specified table space.

TABLESPACE
Names the table space and, optionally, the database it belongs to.

database-name
The name of the database that the table space belongs to.

The **default** is **DSNDB04**.

tablespace-name
The name of the table space.

SHRLEVEL
Tells whether other programs that access the table space while RUNSTATS is running must use read-only access, or can change the table space.

REFERENCE
Allows only read-only access by other programs.

The **default** is **REFERENCE**.

CHANGE
Allows other programs to change the table space or index. With SHRLEVEL CHANGE, uncommitted data can be used to collect statistical summaries.

REPORT
Determines if a set of messages is generated to report the collected statistics.

NO
Indicates that the set of messages is not output to SYSPRINT.

The **default** is **NO**.

YES
Indicates that the set of messages is output to SYSPRINT. The messages generated are dependent on the combination of keywords (such as TABLESPACE, INDEX, TABLE, and COLUMN) specified with the RUNSTATS utility. However, these messages are *not* dependent on the specification of the UPDATE option. REPORT YES always generates a report of SPACE and ACCESSPATH statistics.

UPDATE
Tells whether the collected statistics are inserted into the catalog tables. UPDATE also allows you to select statistics used for access path selection or

statistics used by database administrators. What you specify under the UPDATE option affects what options you can specify under the HISTORY option.

ALL Indicates that all collected statistics will be updated in the catalog.

The **default** is **ALL**.

ACCESSPATH
Indicates that only the catalog table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only the catalog table columns that provide statistics to help the database administrator assess the status of a particular table space or index are updated.

NONE Indicates that no catalog tables are updated with the collected statistics. This option is only valid when REPORT YES is specified.

HISTORY
Records all catalog table inserts or updates to the catalog history tables. What you specify under the UPDATE option affects what options you can specify under the HISTORY option. The table below lists the HISTORY options allowed for each UPDATE option specified.

The default will be supplied by the value specified in STATISTICS HISTORY on panel DSNTIPO.

ALL Indicates that all collected statistics will be updated in the catalog history tables.

ACCESSPATH
Indicates that only the catalog history table columns that provide statistics used for access path selection are updated.

SPACE
Indicates that only SPACE-related catalog statistics are updated in catalog history tables.

NONE Indicates that no catalog history tables are updated with the collected statistics.

Table 84. Allowable HISTORY options based on selected UPDATE option

UPDATE option	HISTORY options
ALL	ALL, ACCESSPATH, SPACE, NONE
ACCESSPATH	ACCESSPATH, NONE
SPACE	SPACE, NONE
NONE	NONE

FORCEROLLUP
Determines if aggregation or rollup of statistics will take place when RUNSTATS is executed in spite of some parts being empty. This will enable the optimizer to select the best access path.

The following are the available options for the FORCEROLLUP Keyword.

YES Indicates that forced aggregation or rollup processing is to be done, even though some parts may not contain data.

NO Indicates that aggregation or rollup will be done only if data is available for all parts.

If data is not available for all parts, DSNU623I message will be issued, provided the installation value for STATISTICS ROLLUP on panel DSNTIPO is set to NO.

KEYCARD
Collects all of the distinct values in all of the 1 to *n* key column combinations for the specified indexes. *n* is the number of columns in the index.

FREQVAL
Controls the collection of frequent value statistics. If you specify FREQVAL, it must be followed by two additional keywords:

NUMCOLS
Indicates the number of key columns to concatenate together when collecting frequent values from the specified index. Specifying '3' means to collect frequent values on the concatenation of the first three key columns. The **default** is **1**, which means collect frequent values on the first key column of the index.

COUNT
Indicates the number of frequent values to be collected. Specifying '15' means collect 15 frequent values from the specified key columns. The **default** is **10**.

Instructions for running RUNSTATS

To run RUNSTATS, you must:

1. Read "Before running RUNSTATS" in this chapter.
2. Prepare the necessary data sets, as described in "Data sets used by RUNSTATS" on page 442.
3. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for RUNSTATS, see "Sample control statements" on page 453.)
4. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 442. (For a complete description of the syntax and options for RUNSTATS, see "Syntax and options of the control statement" on page 432.)
5. Check the compatibility table in "Concurrency and compatibility" on page 444 if you want to run other jobs concurrently on the same target objects.
6. Plan for restart if the RUNSTATS job doesn't complete, as described in "Terminating or restarting RUNSTATS" on page 443. RUNSTATS can be restarted, but it starts over again from the beginning.
7. Run RUNSTATS.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for a description of ways to execute DB2 utilities.

Before running RUNSTATS

The columns RUNSTATS updates can be updated manually using SQL. Use caution when running RUNSTATS after another user has updated the statistical columns of the catalog. Because RUNSTATS puts information in these columns, values changed by the user are replaced.

Data sets used by RUNSTATS

Table 85 describes the data sets used by RUNSTATS. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 85. Data sets used by RUNSTATS

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes

The following objects are named in the utility control statement and do not require DD cards in the JCL:

Table space or index
Object to be scanned. It is named in the RUNSTATS control statement and is accessed through the DB2 catalog.

Creating the control statement

See "Syntax and options of the control statement" on page 432 for RUNSTATS syntax and option descriptions. See "Sample control statements" on page 453 for examples of RUNSTATS usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

"Deciding when to use RUNSTATS"
"Assessing table space status"
"Updating statistics for a partitioned table space"
"Running RUNSTATS on the DB2 catalog" on page 443
"Improving performance" on page 443

Deciding when to use RUNSTATS

DB2 uses the statistics generated by RUNSTATS to determine access paths to data. If no statistics are available, DB2 makes fixed default assumptions. To ensure the effectiveness of the paths selected, use RUNSTATS:

- After a table is loaded
- After an index is physically created
- After a table space is reorganized and inline statistics were not collected
- After there have been extensive updates, deletions, or insertions in a table space
- After you have run RECOVER TABLESPACE, REBUILD INDEX, or REORG INDEX, and you did not collect inline statistics with that utility
- Before running REORG with the OFFPOSLIMIT, INDREFLIMIT, or LEAFDISTLIMIT options.

Assessing table space status

Changes to a table space can also change its space requirements and performance. A database administrator can use RUNSTATS to assess the current status of the table space and help decide whether to reorganize or redesign the table space.

Updating statistics for a partitioned table space

If statistics do not exist for every partition, then RUNSTATS does not compute aggregate statistics (used for access path selection). After newly created partitioned

table spaces have been loaded, run RUNSTATS on the entire table space (or on every partition) to take best advantage of access path selection.

Running RUNSTATS on the DB2 catalog

RUNSTATS may be used for the DB2 catalog, for index space and table space statistics. The following sample execution shows part of the output of RUNSTATS against a catalog table space and its indexes:

```
DSNU000I   DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = DSNTEX
DSNU050I   DSNUGUTC - RUNSTATS TABLESPACE DSNDB06.SYSDBASE INDEX(ALL)
DSNU610I # DSNUSUTP - SYSTABLEPART CATALOG UPDATE FOR DSNDB06.SYSDBASE SUCCESSFUL
DSNU610I # DSNUSUTS - SYSTABLESPACE CATALOG UPDATE FOR DSNDB06.SYSDBASE SUCCESSFUL
DSNU610I # DSNUSUTB - SYSTABLES CATALOG UPDATE FOR SYSIBM.SYSTABLESPACE SUCCESSFUL

DSNU610I # DSNUSUTB - SYSTABLES CATALOG UPDATE FOR SYSIBM.SYSSYNONYMS SUCCESSFUL
DSNU610I # DSNUSUIX - SYSINDEXES CATALOG UPDATE FOR SYSIBM.DSNDSX01 SUCCESSFUL
DSNU610I # DSNUSUIP - SYSINDEXPART CATALOG UPDATE FOR SYSIBM.DSNDSX01 SUCCESSFUL
DSNU610I # DSNUSUCO - SYSCOLUMNS CATALOG UPDATE FOR SYSIBM.DSNDSX01 SUCCESSFUL
DSNU610I # DSNUSUFL - SYSFIELDS CATALOG UPDATE FOR SYSIBM.DSNDSX01 SUCCESSFUL

DSNU610I # DSNUSUIX - SYSINDEXES CATALOG UPDATE FOR SYSIBM.DSNDYX01 SUCCESSFUL
DSNU610I # DSNUSUIP - SYSINDEXPART CATALOG UPDATE FOR SYSIBM.DSNDYX01 SUCCESSFUL
DSNU610I # DSNUSUCO - SYSCOLUMN CATALOG UPDATE FOR SYSIBM.DSNDYX01 SUCCESSFUL
DSNU610I # DSNUSUFL - SYSFIELDS CATALOG UPDATE FOR SYSIBM.DSNDYX01 SUCCESSFUL
DSNU010I   DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

DB2 uses the statistics collected on the catalog to determine the access path for user queries.

Improving performance

You can specify the STATISTICS keyword in LOAD, REBUILD INDEX, and REORG utility statements, which results in updated table space or index space catalog statistics for the objects the utility was run on. Another method of improving RUNSTATS performance is to specify the SAMPLE option on table spaces that were defined with the LARGE option, which reduces the number of rows sampled for statistics.

When you run RUNSTATS concurrently against partitions of a partitioned table space or index, the sum of the processor time for the concurrent jobs will be roughly equivalent to the processor time it takes to run a single RUNSTATS job against the entire table space or index. However, the total elapsed time for the concurrent jobs can be significantly less than when you run RUNSTATS against an entire table space or index.

When requesting nonindexed column statistics, provide a list of columns that might be used in queries as search conditions in a WHERE clause. Collecting statistics on all columns of a table is costly and might not be necessary.

Terminating or restarting RUNSTATS

You can restart a RUNSTATS utility job, but it starts from the beginning again.

You can terminate RUNSTATS with the TERM UTILITY command.

For guidance in restarting online utilities, see “Restarting an online utility” on page 46.

Concurrency and compatibility

Individual data and index partitions are treated as distinct target objects. Utilities operating on different partitions of the same table space or index space are compatible.

Table 86 shows which claim classes RUNSTATS claims and drains and any restrictive state the utility sets on the target object.

Table 87 shows which utilities can run concurrently with RUNSTATS on the same target object. The target object can be a table space, an index space, or a partition of a table space or index space. If compatibility depends on particular options of a utility, that is also shown.

Table 86. Claim classes of RUNSTATS operations. Use of claims and drains; restrictive states set on the target object.

Target	RUNSTATS TABLESPACE SHRLEVEL REFERENCE	RUNSTATS TABLESPACE SHRLEVEL CHANGE	RUNSTATS INDEX SHRLEVEL REFERENCE	RUNSTATS INDEX SHRLEVEL CHANGE
Table space or partition	DW/UTRO	CR/UTRW[1]		
Index or partition			DW/UTRO	CR/UTRW

Legend:
- DW - Drain the write claim class - concurrent access for SQL readers
- CR - Claim the read claim class
- UTRO - Utility restrictive state - read only access allowed
- UTRW - Utility restrictive state - read/write access allowed
- Blank - Object is not affected by this utility.

Notes:

1. If the target object is a segmented table space, SHRLEVEL CHANGE does not allow you to concurrently execute an SQL *searched* DELETE without the WHERE clause.

Table 87. RUNSTATS compatibility

	RUNSTATS TABLESPACE SHRLEVEL REFERENCE	RUNSTATS TABLESPACE SHRLEVEL CHANGE	RUNSTATS INDEX SHRLEVEL REFERENCE	RUNSTATS INDEX SHRLEVEL CHANGE
CHECK DATA DELETE YES	No	No	No	No
CHECK DATA DELETE NO	Yes	Yes	Yes	Yes
CHECK INDEX	Yes	Yes	Yes	Yes
CHECK LOB	Yes	Yes	Yes	Yes
COPY INDEXSPACE	Yes	Yes	Yes	Yes
COPY TABLESPACE	Yes	Yes	Yes	Yes
DIAGNOSE	Yes	Yes	Yes	Yes
LOAD	No	No	No	No
MERGECOPY	Yes	Yes	Yes	Yes
MODIFY RECOVERY	Yes	Yes	Yes	Yes
QUIESCE	Yes	Yes	Yes	Yes

Table 87. RUNSTATS compatibility (continued)

	RUNSTATS TABLESPACE SHRLEVEL REFERENCE	RUNSTATS TABLESPACE SHRLEVEL CHANGE	RUNSTATS INDEX SHRLEVEL REFERENCE	RUNSTATS INDEX SHRLEVEL CHANGE
REBUILD INDEX	Yes	Yes	No	No
RECOVER INDEX	Yes	Yes	No	No
RECOVER TABLESPACE (no options)	No	No	Yes	Yes
RECOVER TABLESPACE TOCOPY or TORBA	No	No	No	No
RECOVER ERROR RANGE	No	No	Yes	Yes
REORG INDEX	Yes	Yes	No	No
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No	No	No	No
REORG TABLESPACE UNLOAD ONLY or EXTERNAL	Yes	Yes	Yes	Yes
REPAIR DUMP or VERIFY	Yes	Yes	Yes	Yes
REPAIR LOCATE KEY or RID DELETE or REPLACE	No	No	No	Yes
REPAIR LOCATE TABLESPACE PAGE REPLACE	No	No	Yes	Yes
REPAIR LOCATE INDEX PAGE REPLACE	Yes	Yes	No	No
REPORT	Yes	Yes	Yes	Yes
RUNSTATS	Yes	Yes	Yes	Yes
STOSPACE	Yes	Yes	Yes	Yes
UNLOAD	Yes	Yes	Yes	Yes

Reviewing RUNSTATS output

RUNSTATS alters the tables and columns in the DB2 catalog tables listed below. A report of statistics gathered during processing is generated with the REPORT YES option.

RUNSTATS sets the following columns to -1 for large table spaces.
- CARD in SYSTABLES
- CARD in SYSINDEXPART
- FAROFFPOS in SYSINDEXPART
- NEAROFFPOS in SYSINDEXPART
- FIRSTKEYCARD in SYSINDEXES
- FULLKEYCARD in SYSINDEXES

Index statistics and table space statistics: The following catalog tables are updated depending on the source of the statistics as well as the value of the UPDATE option.

RUNSTATS

Table 88. Catalog tables updated by RUNSTATS

Keyword	UPDATE Option	Catalog Table with HISTORY ALL[4]	Catalog Table with HISTORY ACCESSPATH	Catalog Table with HISTORY SPACE
TABLESPACE	UPDATE ALL	SYSTABLESPACE SYSTABLEPART[2] SYSTABLES[2] SYSTABSTATS[1,2] SYSLOBSTATS[3]	SYSTABLESPACE SYSTABLES[2] SYSTABSTATS[1,2]	SYSTABLEPART[2] SYSLOBSTATS[3]
	UPDATE ACCESSPATH[2]	SYSTABLESPACE SYSTABLES SYSTABSTATS[1]	SYSTABLESPACE SYSTABLES SYSTABSTATS[1]	
	UPDATE SPACE[2]	SYSTABLEPART SYSLOBSTATS SYSTABLES		SYSTABLEPART SYSLOBSTATS SYSTABLES
TABLE	UPDATE ALL	SYSCOLUMNS SYSCOLSTATS[1]	SYSCOLUMNS SYSCOLSTATS[1]	
	UPDATE ACCESSPATH	SYSCOLUMNS SYSCOLSTATS[1]	SYSCOLUMNS SYSCOLSTATS[1]	
INDEX	UPDATE ALL	SYSCOLUMNS SYSCOLDIST SYSCOLDISTSTATS[1] SYSCOLSTATS[1] SYSINDEXES SYSINDEXPART SYSINDEXSTATS[1]	SYSCOLUMNS SYSCOLDIST SYSCOLDISTSTATS[1] SYSCOLSTATS[1] SYSINDEXPART SYSINDEXSTATS[1]	SYSINDEXES
	UPDATE ACCESSPATH	SYSCOLUMNS SYSCOLDIST SYSCOLDISTSTATS[1] SYSCOLSTATS SYSINDEXES SYSINDEXSTATS[1]	SYSCOLUMNS SYSCOLDIST SYSCOLDISTSTATS[1] SYSCOLSTATS SYSINDEXES SYSINDEXSTATS[1]	SYSINDEXES
	UPDATE SPACE	SYSINDEXPART		SYSINDEXPART

Notes:

1. Only updated for partitioned cases. When you run RUNSTATS against single partitions of an object, the partition-level statistics that result are used to update the aggregate statistics for the entire object. The catalog tables containing these partition-level statistics are the following:
   - SYSCOLSTATS
   - SYSCOLDISTSTATS
   - SYSTABSTATS
   - SYSINDEXSTATS
2. If the specified table space is a LOB table space, this is not applicable.
3. Only applicable when the specified table space is a LOB table space.
4. When HISTORY NONE is specified, none of the catalog tables will be updated.

Access path statistics

The catalog table columns listed in Table 89 are used by DB2 to select access paths to data during the bind process. Refer to Part 5 (Volume 2) of *DB2 Administration Guide* for further information regarding these columns.

Table 89 does not describe information about LOB columns, because those statistics are not used for access path selection. For indexes on auxiliary tables, only the NLEVELS and FIRSTKEYCARDF columns in SYSIBM.SYSINDEXES have an effect on the access path. For information on what values in these columns indicate for LOBs, see Appendix D of *DB2 SQL Reference*.

A value in the column "Use" indicates whether information about the DB2 catalog column is General-use Programming Interface and Associated Guidance Information (G) or Product-sensitive Programming Interface and Associated Guidance Information (S), as defined in "Notices" on page 725.

Table 89. DB2 catalog columns used to select data access paths

Column Name	Column Description	Use
SYSTABLES		
CARDF	Total number of rows in the table.	S
NPAGES	Total number of pages on which rows of this table appear.	S
NPAGESF	Total number of pages used by the table.	S
PCTROWCOMP	Percentage of rows compressed within the total number of active rows in the table.	S
STATSTIME	The date and time when RUNSTATS was last executed to update the statistics.	G
SYSTABSTATS		
CARD or CARDF	Total number of rows in the partition.	S
NPAGES	Total number of pages on which rows of this partition appear.	S
SYSCOLUMNS		
COLCARDF	Estimated number of distinct values for the column. For an indicator column, this is the number of non-zero length, non-null LOBs. The value is -1 if statistics have not been gathered. The value is -2 for columns of an auxiliary table.	S
HIGH2KEY	Second highest value of the column. Blank if statistics have not been gathered or the column is an indicator column or a column of an auxiliary table. If the column has a non-character data type, the data might not be printable. This is an updatable column.	S
LOW2KEY	Second lowest value of the column. Blank if statistics have not been gathered or the column is an indicator column or a column of an auxiliary table. If the column has a non-character data type, the data might not be printable. This is an updatable column.	S
STATSTIME	The date and time when RUNSTATS was last executed to update the statistics.	G

Table 89. DB2 catalog columns used to select data access paths (continued)

Column Name	Column Description	Use
SYSCOLDIST		
CARDF	The number of distinct values for the column group. This number is only valid for cardinality (TYPE **C**) key column statistics.	S
COLGROUPCOLNO	Identifies the set of columns associated with the key column statistics.	S
COLVALUE	Actual index column value that is being counted for distribution index statistics.	S
FREQUENCYF	Percentage of rows, multiplied by 100, that contain the values specified in COLVALUE.	S
NUMCOLUMNS	The number of columns associated with the key column statistics	G
STATSTIME	The date and time when RUNSTATS was last executed to update the statistics.	G
SYSTABLESPACE		
NACTIVE or NACTIVEF	Number of active pages in the table space; shows the number of pages that are touched if a record cursor is used to scan the entire file. The value is -1 if statistics have not been gathered.	S
STATSTIME	The date and time when RUNSTATS was last executed to update the statistics.	G
DSSIZE	Maximum size of a data set in kilobytes.	G
SYSINDEXES		
CLUSTERRATIOF	A number between 0 and 1 that when multiplied by 100 gives the percentage of rows in clustering order. For example, a value of 1 indicates that all rows are in clustering order. A value of .87825 indicates that 87.825% rows are in clustering order.	S
CLUSTERING	Whether CLUSTER was specified when the index was created.	G
FIRSTKEYCARDF	Number of distinct values of the first key column.	S
FULLKEYCARDF	Number of distinct values of the full key.	S
NLEAF	Number of leaf pages in the index.	S
NLEVELS	Number of levels in the index tree.	S
STATSTIME	The date and time when RUNSTATS was last executed to update the statistics.	G
SYSINDEXSTATS		
CLUSTERRATIOF	A number between 0 and 1 that when multiplied by 100 gives the percentage of rows in clustering order. For example, a value of 1 indicates that all rows are in clustering order. A value of .87825 indicates that 87.825% rows are in clustering order.	S

Table 89. DB2 catalog columns used to select data access paths (continued)

Column Name	Column Description	Use
FIRSTKEYCARD or FIRSTKEYCARDF	For the index partition, number of distinct values of the first key column.	S
FULLKEYCARD or FULLKEYCARDF	For the index partition, number of distinct values of the key.	S
NLEAF	Number of active leaf pages in the index partition.	S
NLEVELS	Number of levels in the index tree in the partition.	S
KEYCOUNT or KEYCOUNTF	Total number of rows in the partition.	S

Space statistics (columns for tuning information)

The following catalog table columns are updated by RUNSTATS to help database administrators assess the status of a particular table space or index.

A value in the column "Use" indicates whether information about the DB2 catalog column is General-use Programming Interface and Associated Guidance Information (G) or Product-sensitive Programming Interface and Associated Guidance Information (S), as defined in "Notices" on page 725.

Table 90. DB2 catalog columns for tuning information

Column Name	Column Description	Use
SYSTABLEPART		
CARD	Total number of rows in the table space or partition, or number of LOBs in the table space if the table space is a LOB table space. The value is -1 if statistics have not been gathered. The database administrator can validate design assumptions against this actual count. Over a period of time, it can show the rate of change or growth of the table space.	G
DSNUM	Number of datasets.	G
EXTENTS	Number of dataset extents.	G
NEARINDREF	Number of rows relocated near their original page. See the description following FARINDREF.	S

Table 90. DB2 catalog columns for tuning information (continued)

Column Name	Column Description	Use
FARINDREF	Number of rows relocated far from their original page. If an update operation increases the length of a record by more than the amount of space available in the page in which it is stored, the record is moved to another page. Until the table space is reorganized, the record requires an additional page reference when it is accessed. The sum of NEARINDREF and FARINDREF is the total number of such records. For nonsegmented table spaces, a page is considered “near” the present page if the two page numbers differ by 16 or less; otherwise, it is “far from” the present page. For segmented table spaces, a page is considered “near” the present page if the two page numbers differ by (SEGSIZE * 2) or less. Otherwise, it is “far from” its original page. A record relocated near its original page tends to be accessed more quickly than one relocated far from its original page.	S
PAGESAVE	Percentage of pages saved in the table space or partition as a result of using data compression. For example, a value of 25 indicates a savings of 25 percent, so that the pages required are only 75 percent of what would be required without data compression. The value is 0 if there are no savings from using data compression, or if statistics have not been gathered. The value can be negative if using data compression causes an increase in the number of pages in the data set. This calculation includes the overhead bytes for each row, the bytes required for the dictionary, and the bytes required for the current FREEPAGE and PCTFREE specification for the table space and partition. This calculation is based on an average row length and the result varies depending on the actual lengths of the rows.	S
PERCACTIVE	Percentage of space occupied by rows of data from active tables. The value is -1 if statistics have not been gathered. The value is -2 if the table space is a LOB table space. A database administrator can use this figure to validate design assumptions, and tell how much of the space allocated to the table space is utilized. This value is influenced by the PCTFREE and the FREEPAGE parameters on the CREATE TABLESPACE statement, and by unused segments of segmented table spaces.	S

Table 90. DB2 catalog columns for tuning information (continued)

Column Name	Column Description	Use
PERCDROP	For nonsegmented table spaces, the percentage of space occupied by rows of data from dropped tables. For segmented table spaces, this value is zero. After reorganization, this value is always zero. Space occupied by dropped tables is reclaimed by reorganization. Hence, this figure is one indicator of when a table space should be reorganized.	S
CARDF	Total number of rows in the table space or partition, or if the table space is a LOB table space, the number of LOBs in the table space. The value is -1 if statistics have not been gathered. The database administrator can validate design assumptions against this actual count. Over a period of time, it can show the rate of change or growth of the table space.	G
SPACE	The number of kilobytes of space currently allocated for all extents. A value of -1 indicates that the data set was defined with the DEFINE NO attribute, and the first insert operation has not occurred.	G
SPACEF	The number of kilobytes of DASD storage.	G
PQTY (user-managed)	The primary space allocation in 4-KB blocks for the data set.	G
SQTY (user-managed)	The secondary space allocation in 4-KB blocks for the data set, in small integer format.	G
SECQTYI (user-managed)	The secondary space allocation in 4-KB blocks for the data set, in integer format.	G
SYSINDEXPART		
CARDF	Number of rows referred to by the index or partition. Those figures, for all the partitions, tell the database administrator how the key ranges specified for each partition have divided the rows among the several partitions.	S
DSNUM	Number of datasets.	G
EXTENTS	Number of dataset extents.	G

Table 90. DB2 catalog columns for tuning information (continued)

Column Name	Column Description	Use
FAROFFPOSF	Number of times is would be necessary to access a different, “far-off” page when accessing all the data records in index order. Each time, it is almost certain that accessing the “next” record would require I/O activity. For nonsegmented table spaces, a page is considered far-off the present page if the two page numbers differ by 16 or more. For segmented table spaces, a page is considered far-off the present page if the two page numbers differ by SEGSIZE * 2 or more. Together, NEAROFFPOS and FAROFFPOS tell how well the index follows the cluster pattern of the table space. For a clustering index, NEAROFFPOS and FAROFFPOS approach a value of 0 as clustering improves. A reorganization should bring them nearer their optimal values; however, if a nonzero FREEPAGE value was specified on the CREATE TABLESPACE statement, the NEAROFFPOS after reorganization reflects the table on which the index is defined. Optimal values should not be expected for nonclustering indexes. Not applicable for the index on an auxiliary table (-1).	S
LEAFNEAR	Number of leaf pages located physically near previous leaf pages for successive active leaf pages.	S
LEAFFAR	Number of leaf pages located physically far away from previous leaf pages for successive active leaf pages accessed in an index scan.	S
NEAROFFPOSF	Number of times it would be necessary to access a different, “near-off” page when accessing all the data records in index order. Each time, it is probable that accessing the “next” record would require I/O activity. See the description following FAROFFPOS. NEAROFFPOS is incremented if the current indexed row is not on the same or next data page of the previous indexed row, and the distance between the two data pages does not qualify for FAROFFPOS. For nonsegmented table spaces, a page is considered near-off the present page if the difference between the two page numbers is greater than or equal to 2, and less than 16. For segmented table spaces, a page is considered near-off the present page if the difference between the two page numbers is greater than or equal to 2, and less than SEGSIZE * 2. A nonzero value in the NEAROFFPOS field after a REORG might be attributed to the number of space map pages contained in the segmented table space.Not applicable for the index on an auxiliary table (-1).	S

Table 90. DB2 catalog columns for tuning information (continued)

Column Name	Column Description	Use
LEAFDIST	100 times the average distance in page IDs between successive leaf pages during a sequential access of the index. This value helps to tell how well an index is organized. The value is at its lowest just after the index has been reorganized. Changes increase it; and you can reduce it again by reorganizing the index, either explicitly or as part of a general table space reorganization.	S
PSEUDO_DEL_ENTRIES	Number of pseudo deleted keys.	S
SPACE	The number of kilobytes of space currently allocated for all extents. A value of -1 indicates that the data set was defined with the DEFINE NO attribute, and the first insert operation has not occurred.	G
SPACEF	The number of kilobytes of DASD storage.	G
PQTY (user-managed)	The primary space allocation in 4-KB blocks for the data set.	G
SQTY (user-managed)	The secondary space allocation in 4-KB blocks for the data set, in small integer format.	G
SECQTYI (user-managed)	The secondary space allocation in 4-KB blocks for the data set, in integer format.	G
SYSLOBSTATS		
FREESPACE	The number of kilobytes of available space in the LOB table space up to the high-used RBA.	S
ORGRATIO	The ratio of organization in the LOB table space. A value of 1 indicates perfect organization of the LOB table space. The greater the value exceeds 1, the more disorganized the LOB table space.	S

After running RUNSTATS

After running RUNSTATS, rebind any application plans that use the tables or indexes so that they use the new statistics.

Sample control statements

Example 1: Update catalog statistics while allowing changes. Update the catalog statistic columns for table space DSN8S71E and all its associated indexes, sampling 25 percent of the rows. Permit other processes to make changes while this utility is executing.

```
//STEP1    EXEC DSNUPROC,UID='IUJQU225.RUNSTA',TIME=1440,
//         UTPROC='',
//         SYSTEM='V71A',DB2LEV=DB2A
//UTPRINT  DD  SYSOUT=*
//SYSIN    DD *
```

```
RUNSTATS TABLESPACE DSN8D71A.DSN8S71E
   TABLE(ALL) SAMPLE 25
    INDEX(ALL)
   SHRLEVEL CHANGE
```

Example 2: Update catalog statistics, do not allow updates. Update the catalog statistics for indexes XEMPL1 and XEMPL2. Do not permit other processes to change the table space associated with XEMPL1 and XEMPL2 (table space DSN8S71E) while this utility is executing.

```
RUNSTATS INDEX (DSN8710.XEMPL1,DSN8710.XEMPL2)
```

Example 3: Update index statistics. Obtain statistics on the index XEMPL1.

```
RUNSTATS INDEX (DSN8710.XEMPL1)
```

Example 4: Update statistics for several tables. Update the catalog statistics for all columns in the TCONA and TOPTVAL tables in table space DSN8D71P.DSN8S71C. Update the column statistics for the LINENO and DSPLINE columns in the TDSPTXT table in table space DSN8D71P.DSN8S71C.

```
RUNSTATS TABLESPACE(DSN8D71P.DSN8S71C) TABLE (TCONA)
                                       TABLE (TOPTVAL) COLUMN(ALL)
                                       TABLE (TDSPTXT) COLUMN(LINENO,DSPLINE)
```

Example 5: Update all statistics for a table space. Update all catalog statistics (table space, tables, columns, and indexes) for a table space.

```
RUNSTATS TABLESPACE(DSN8D71P.DSN8S71C) TABLE INDEX
```

Example 6: Update statistics used for access path selection. Update the catalog with *only* the statistics that are collected for access path selection. Report the collected statistics and route the statistics to SYSPRINT.

```
RUNSTATS TABLESPACE DSN8D71A.DSN8S71E
  REPORT YES
  UPDATE ACCESSPATH
```

Example 7: Update all statistics and generate report. Update the catalog with *all* the statistics (access path and space). Report the collected statistics and route the statistics to SYSPRINT.

```
RUNSTATS TABLESPACE DSN8D71A.DSN8S71E
  REPORT YES
  UPDATE ALL
```

Example 8: Report statistics without updating catalog. Do not update the catalog with the collected statistics. Report the collected statistics and route the statistics to SYSPRINT.

```
RUNSTATS TABLESPACE DSN8D71A.DSN8S71E
  REPORT YES
  UPDATE NONE
```

Example 9: Update statistics for a partition. Update the statistics for the table space and the partitioning index after a change to partition 1.

```
RUNSTATS TABLESPACE DSN8D71A.DSN8S71E PART 1 INDEX(DSN8710.XEMP1 PART 1)
```

Example 10: Report statistics and update catalog and history tables. Update the catalog with all statistics and update the history tables with all catalog table inserts or updates. Report the collected statistics and route the statistics to SYSPRINT.

```
RUNSTATS TABLESPACE DBOE0101.TLOE0101
           INDEX
           TABLE
           SHRLEVEL REFERENCE
           REPORT YES
           UPDATE ALL
           HISTORY ALL
```

Example 11: Update Keycard and Freqvals Statistics: Update the statistics on the index IXNPI. Collect all the distinct values in all the key column combinations and control the collection of frequent value statistics.

```
RUNSTATS INDEX (SYSADM.IXNPI)
  KEYCARD
  FREQVAL NUMCOLS 1 COUNT 5
  FREQVAL NUMCOLS 2 COUNT 10
 REPORT YES
```

Chapter 27. STOSPACE

The STOSPACE online utility updates DB2 catalog columns that indicate how much space is allocated for storage groups and related table spaces and indexes.

How to use this chapter: For a diagram of STOSPACE syntax and a description of available options, see "Syntax and options of the control statement". For detailed guidance on running this utility, see "Instructions for running STOSPACE" on page 458.

Output: The output from STOSPACE consists of new values in a number of catalog tables. See "Reviewing STOSPACE output" on page 461 for a list of columns and tables STOSPACE updates.

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- STOSPACE privilege
- SYSCTRL or SYSADM authority

Execution phases of STOSPACE: The STOSPACE utility operates in these phases:

Phase Description
UTILINIT
Initialization and setup
STOSPACE
UTILTERM
Cleanup

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

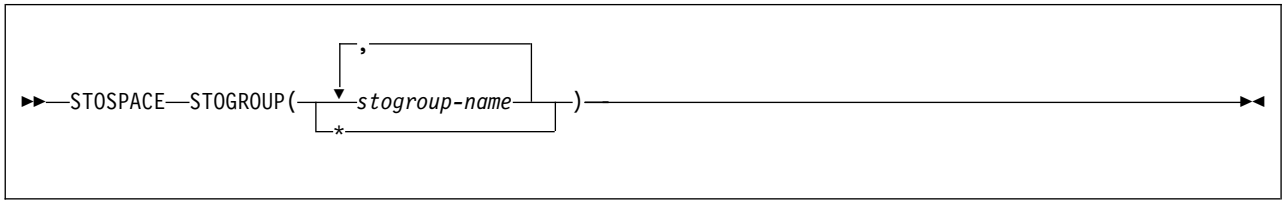

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

© Copyright IBM Corp. 1983, 2001

STOGROUP
Identifies the groups to be processed.

(*stogroup-name***, ...)**
Is the name of a storage group. You can use a list of from one to eight storage group names. Separate items in the list by commas and enclose them in parentheses.

***** Processes all storage groups.

Instructions for running STOSPACE

To run STOSPACE, you must:

1. Prepare the necessary data sets, as described in "Data sets used by STOSPACE".
2. Create JCL statements, by using one of the methods described in "Chapter 3. Invoking DB2 online utilities" on page 23. (For examples of JCL for STOSPACE, see "Sample control statement" on page 461.)
3. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in "Instructions for specific tasks" on page 459. (For a complete description of the syntax and options for STOSPACE, see "Syntax and options of the control statement" on page 457.)
4. Check the compatibility rules in "Concurrency and compatibility" on page 460 if you want to run other jobs concurrently on the same target objects.
5. Plan for restart if the STOSPACE job doesn't complete, as described in "Terminating or restarting STOSPACE" on page 460.
6. Run STOSPACE.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Data sets used by STOSPACE

Table 91 describes the data sets used by STOSPACE. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 91. Data sets used by STOSPACE

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes

The following object is named in the utility control statement and does not require a DD card in the JCL:

Storage group
Object to be reported. It is named in the STOSPACE control statement and is accessed through the DB2 catalog.

Creating the control statement

See "Syntax and options of the control statement" on page 457 for STOSPACE syntax and option descriptions. See "Sample control statement" on page 461 for examples of STOSPACE usage.

Instructions for specific tasks

To perform the following tasks, specify the options and values for those tasks in your utility control statement:

"Ensuring availability of objects required by STOSPACE"
"Obtaining statistical information with STOSPACE"
"Understanding the values in a SPACE column" on page 460

Ensuring availability of objects required by STOSPACE

For each specified storage group, STOSPACE looks at the SYSIBM.SYSTABLESPACE and SYSIBM.SYSINDEXES catalog tables to tell which objects belong to that storage group. For each object, the amount of space allocated is determined from an appropriate VSAM catalog. Hence the table spaces and indexes need not be available to DB2 when STOSPACE is running; only the DB2 catalog and appropriate VSAM catalogs are required. However, to gain access to the VSAM catalog, the utility must have available to it the DBD for the objects involved. This requires that the appropriate database, table spaces, and index spaces not be in the stopped state.

Obtaining statistical information with STOSPACE

Table 92 lists statistical information recorded by the STOSPACE utility that is useful for space allocation decisions.

Table 92. DB2 catalog data collected by STOSPACE

Catalog Table	Column Name	Column Description
SYSTABLESPACE	SPACE	Number of kilobytes of storage allocated to the table space
SYSTABLEPART	SPACE	Number of kilobytes of storage allocated to the table space partition
SYSINDEXES	SPACE	Number of kilobytes of storage allocated to the index
SYSINDEXPART	SPACE	Number of kilobytes of storage allocated to the index partition
SYSSTOGROUP	SPACE	Number of kilobytes of storage allocated to the storage group
SYSSTOGROUP	SPCDATE	Date when STOSPACE was last run on a particular storage group
SYSSTOGROUP	STATSTIME	Time when STOSPACE was last run on a particular storage group

When DB2 storage groups are used in the creation of table spaces and indexes, DB2 defines the data sets for them. The STOSPACE utility permits a site to monitor the DASD space allocated for the storage group.

STOSPACE does not accumulate information for more than one storage group. If a partitioned table space or index space has partitions in more than one storage group, the information in the catalog about that space comes from only the group for which STOSPACE was run.

When you run the STOSPACE utility, the SPACE column of the catalog represents the high allocated RBA of the VSAM linear data set. Use the value in the SPACE column to project space requirements for table spaces, table space partitions, index spaces, and index space partitions over time. Use the output from the access method services LISTCAT command to determine which table spaces and index

spaces have allocated secondary extents; when you find these, it is a good idea to increase the primary quantity value for the data set and run the REORG utility.

For information about space utilization in the DSN8S71E table space in the DSN8D71A database, first run the STOSPACE utility, and then execute this SQL statement:

General-use Programming Interface

```
SELECT SPACE
  FROM SYSIBM.SYSTABLESPACE
  WHERE NAME = 'DSN8S71E
  AND DBNAME = 'DSN8D71A';
```

End of General-use Programming Interface

Alternatively, you can use TSO to look at data set and pack descriptions.

To update SYSIBM.SYSSTOGROUP for storage group DSN8G710, as well as SYSIBM.SYSTABLESPACE and SYSIBM.SYSINDEXES for every table space and index belonging to DSN8G710, use the following utility:

```
STOSPACE STOGROUP DSN8G710
```

Understanding the values in a SPACE column

The value in a SPACE column is total allocated space, not only space allocated on the current list of volumes in the storage groups. Volumes can be deleted from a storage group even though space on those volumes is still allocated to DB2 table spaces or indexes. Deletion of a volume from a storage group prevents future allocations; it does not withdraw a current allocation.

For each specified storage group, STOSPACE looks at the SYSIBM.SYSTABLESPACE and SYSIBM.SYSINDEXES catalog tables to tell which objects belong to that storage group. For each object, the amount of space allocated is determined from an appropriate VSAM catalog. Therefore, the table spaces and indexes need not be available to DB2 when STOSPACE is running; only the DB2 catalog and appropriate VSAM catalogs are needed. However, to gain access to the VSAM catalog, the utility must have available to it the DBD for the objects involved. This requires that the appropriate database, table spaces, and index spaces not be in the stopped state.

Considerations for running STOSPACE

For user-defined spaces, STOSPACE does not record any statistics.

Terminating or restarting STOSPACE

You can terminate STOSPACE with the TERM UTILITY command.

You can restart a STOSPACE utility job; however, it starts again from the beginning.

For more guidance in restarting online utilities, see “Restarting an online utility” on page 46.

Concurrency and compatibility

STOSPACE does not set a utility restrictive state on the target object.

STOSPACE can run concurrently on the same target object with any utility. However, because STOSPACE updates the catalog, concurrent STOSPACE utility jobs or other concurrent applications that update the catalog might cause timeouts and deadlocks.

Reviewing STOSPACE output

The output from STOSPACE consists of new values in the columns and tables listed. In each case, an amount of space is given in kilobytes.

- SPACE in SYSIBM.SYSINDEXES shows the amount of space allocated to indexes. If the index is not defined using STOGROUP, or STOSPACE has not been executed, the value is zero.
- SPACE in SYSIBM.SYSTABLESPACE shows the amount of space allocated to table spaces. If the table space is not defined using STOGROUP, or STOSPACE has not been executed, the value is zero.
- SPACE in SYSIBM.SYSINDEXPART shows the amount of space allocated to index partitions. If the partition is not defined using STOGROUP, or STOSPACE has not been executed, the value is zero.
- SPACE in SYSIBM.SYSTABLEPART shows the amount of space allocated to table partitions. If the partition is not defined using STOGROUP, or STOSPACE has not been executed, the value is zero.
- SPACE in SYSIBM.SYSSTOGROUP shows the amount of space allocated to storage groups.
- SPCDATE in SYSIBM.SYSSTOGROUP shows, in the form *yyddd*, the date when STOSPACE was last used on a particular storage group.
- STATSTIME in SYSIBM.SYSSTOGROUP shows the timestamp for the time.

Sample control statement

Example: Update catalog SPACE columns. Update the DB2 catalog SPACE columns for storage group DSN8G710.

```
//STEP1    EXEC DSNUPROC,UID='FUAUU330.STOSPCE',
//         UTPROC='',
//         SYSTEM='V71A'
//SYSIN    DD *
STOSPACE STOGROUP DSN8G710
//*
```

Chapter 28. TEMPLATE

The TEMPLATE utility control statement lets you allocate data sets, without using JCL DD cards, during the processing of LISTDEF list. In its simplest form, the TEMPLATE control statement defines the data set naming convention. The control statement can optionally contain allocation parameters that define data set size, location, and attributes.

The TEMPLATE control statement uses the MVS DYNALLOC macro (SVC 99) to perform data set allocation. Therefore, the facility is constrained by the limitations of this macro and by the subset of DYNALLOC, supported by TEMPLATE. See the OS/390 MVS Programming: Assembler Services Guide for more details.

Output: The TEMPLATE control statement generates a dynamic allocation template with an assigned name for later reference.

Authorization required: No privileges are required to execute this control statement. When a TEMPLATE is referenced by a specific utility, privileges will be checked at that time.

Execution phases of TEMPLATE: The TEMPLATE control statement executes entirely in the UTILINIT phase, performing setup for the subsequent utility.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

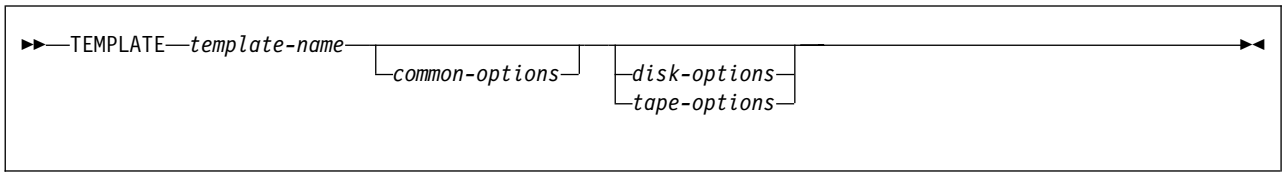

© Copyright IBM Corp. 1983, 2001

TEMPLATE

common-options:

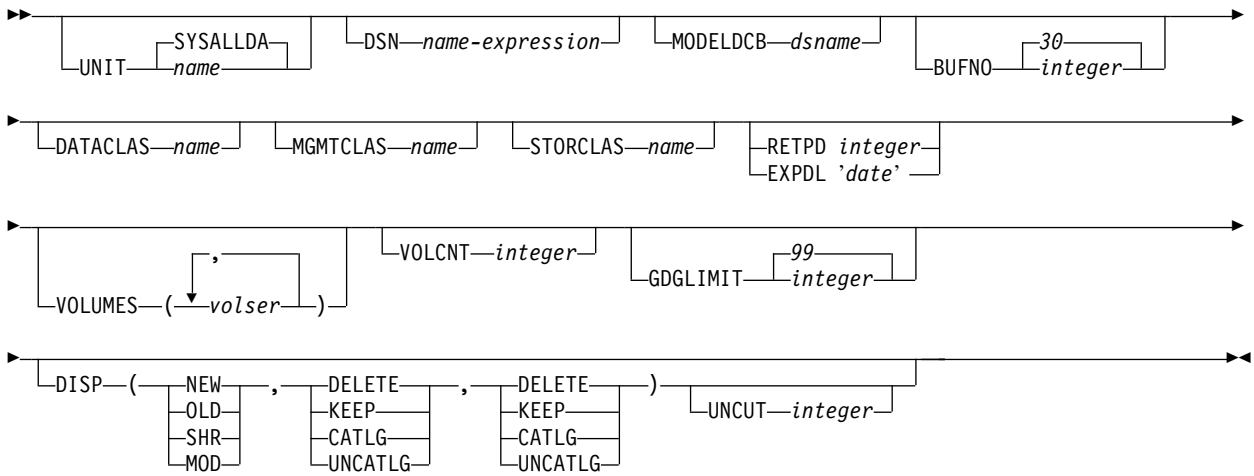

disk-options:

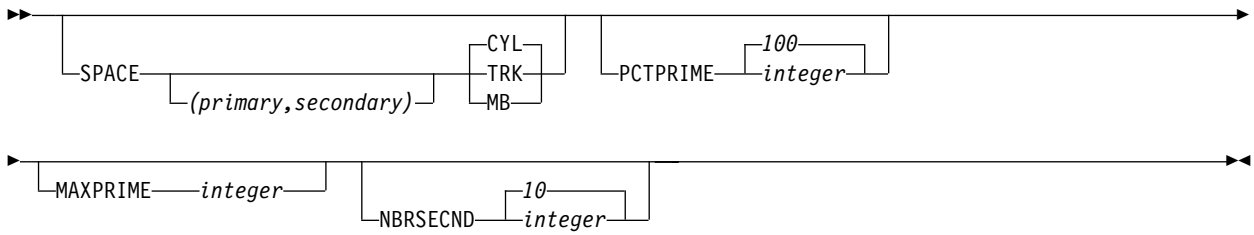

tape-options:

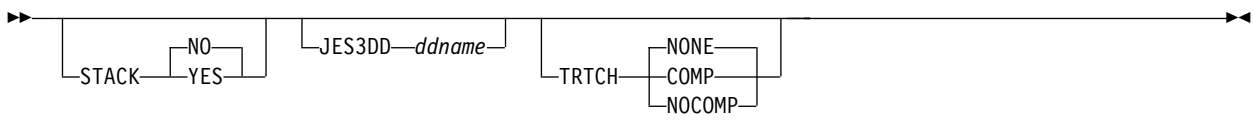

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see "Control statement coding rules" on page 23.

TEMPLATE *template-name*

Defines a data set allocation template and assigns the template a name, *template-name*, for subsequent reference on a DB2 utility

control statement. The *template-name* can have up to eight alphanumeric characters and must begin with an alphabetical character.

The *template-name* is followed by keywords that control the allocation of tape and disk data sets. Disk options and tape options must not appear on the same TEMPLATE statement. The UNIT keyword specifies a generic unit name that is defined on your system. This value is used to determine if a disk or tape data set is being allocated. All other keywords specified on the TEMPLATE control statement must be consistent with the unit type specified.

common-options

This section describes options that apply to both disk and tape data sets.

UNIT Specifies the device-number, device-type (generic) or group-name for the data set. All other TEMPLATE keywords are validated based on the type of unit specified (DASD or TAPE). The **default** is **SYSALLDA**.

DSN *name-expression*

Specifies the template for the MVS data set name. You can specify data set name by using symbolic variables, non-variable alphanumeric or national characters, or any combination of these characters. Each symbolic variable is substituted by its related value at execution time to form a specific data set name. When used in a DSN expression, substitution variables begin with an ampersand sign (&) and end with a period (.). For example:

```
DSN &DB..&TS..D&JDATE..COPY&ICTYPE.
```

See Table 93, Table 95 on page 466, Table 94 on page 466 and Table 96 on page 467 for a list of substitution variables.

If the DSN operand contains the parentheses special characters it must be enclosed in single quotes. For example:

```
TEMPLATE X DSN 'A.GDG.VERSION(+1)'
```

Using &DATE. alone, generates an invalid data set qualifier for all numeric-type variables (all date or time-type variables, and others such as `&SEQ.` or `&PART.`). Examples of valid specification are `P&PART.` and `D&DATE..`

Some substitution variables are invalid if you use TEMPLATE with an incompatible utility. For example, ICTYPE would not be meaningful if the TEMPLATE were used with LOAD SYSDISC. Other variables take on default values when their values are not known. For example, &PART. becomes 00000 for non-partitioned objects.

Parentheses for the DSN name-expression are optional. For example:

```
DSN(&DB..&TS..D&DATE.)
```

Table 93. JOB variables

Variable	Description
`&JOBNAME.` or `&JO.`	The MVS job name.

Table 93. JOB variables (continued)

&STEPNAME. or &ST.	The MVS step name. This variable may be needed if data set names from two different job steps conflict.
&UTILID. or &UT.	The utilid truncated to 8 characters and error checked for invalid DSN characters.
&SSID. or &SS.	Subsystem ID (non-data sharing) or group attach name (data sharing).

Table 94. UTILITY variables

Variable	Description
&ICTYPE. or &IC.	Single character image copy type. Valid only for image copy templates. The substitution is governed by whether a full image copy (F), an incremental image copy (I) or a CHANGELIMIT image copy (C) is specified by user.
&UTILNAME. or &UN.	Special values are assigned to some utilities: CHECKD for CHECK DATA, CHECKI for CHECK INDEX, CHECKL for CHECK LOB, REORGI for REORG INDEX, and REORGT for REORG TABLESPACE. Utility names that are longer than 8 characters are truncated to 8 characters.
&SEQ. or &SQ.	Sequence number of the list item in the list.
&LOCREM. or &LR.	Single character L is used when the utility defines a COPYDDN *ddname1*. The single character R is used when the utility defines a RECOVERYDDN *ddname1*. You can replicate the SYSCOPY ICBACKUP column information by using both the &LOCREM. and &PRIBAC. variables. Valid only for image copy templates.
&PRIBAC. or &PB.	Single character P is used when the utility defines a COPYDDN *ddname2*. The single character B is used when the utility defines a RECOVERYDDN *ddname2*. You can replicate the SYSCOPY ICBACKUP column information by using both the &LOCREM. and &PRIBAC. variables. Valid only for image copy templates.

Table 95. OBJECT variables

Variable	Description
&LIST. or &LI.	The name of the list defined using the LISTDEF control statement and referenced on the same control statement as this template. For use with COPY FILTERDDN templates. All objects in the list are copied to one data set, which makes &TS. and &IS. meaningless.
&DB.	Database name
&TS.	Table space name
&IS.	Index space name
&SN.	Space name (table space or index space)
&PART. or &PA.	Five-digit partition number padded with leading zeros

Table 95. OBJECT variables (continued)

Note: If you use PARTLEVEL keyword, the template should take it into consideration and use this variable. Without this variable, the List Manager builds the list with different part numbers, but the templates generate the same data set name, which results in a duplicate data set name.

Table 96. DATE and TIME variables

Variable	Description
&DATE. or &DT.	YYYYDDD
&TIME. or &TI.	HHMMSS
&JDATE. or &JU.	YYYYDDD
&YEAR. or &YE.	YYYY portion of &DATE.
&MONTH. or &MO.	MM
&DAY. or &DA.	DD
&JDAY. or &JD.	DDD portion of &DATE.
&HOUR. or &HO.	HH portion of &TIME.
&MINUTE. or &MI.	MM portion of &TIME.
&SECOND. or &SC.	SS portion of &TIME.

Note: All date/time values are set using the STCK instruction and reflect the date and time value in Greenwich Mean Time (GMT).

DISP *integer* Specifies the data set disposition by using three positional parameters: status, normal-termination, and abnormal-termination. All three parameters must be specified.

status
Standard MVS values are allowed: NEW, OLD, SHR, MOD.

normal-termination
Standard MVS values are allowed: DELETE, KEEP, CATLG, UNCATLG.

abnormal-termination
Standard MVS values are allowed: DELETE, KEEP, CATLG, UNCATLG.

Default values for DISP vary depending on the utility and the data set being allocated. Defaults for restarted utilities also differ from default values for new utility executions. Default values are shown in Table 97 and Table 98 on page 468.

Table 97. Data dispositions for dynamically allocated data sets

ddname	CHECK DATA	CHECK INDEX or CHECK LOB	COPY	COPY-TOCOPY	LOAD	MERGE-COPY	REBUILD INDEX	REORG INDEX	REORG TABLE-SPACE	UNLOAD
SYSREC	Ignored	Ignored	Ignored	Ignored	OLD KEEP KEEP	Ignored	Ignored	Ignored	NEW CATLG CATLG	NEW CATLG CATLG
SYSDISC	Ignored	Ignored	Ignored	Ignored	NEW CATLG CATLG	Ignored	Ignored	Ignored	NEW CATLG CATLG	Ignored

Table 97. Data dispositions for dynamically allocated data sets (continued)

ddname	CHECK DATA	CHECK INDEX or CHECK LOB	COPY	COPY-TOCOPY	LOAD	MERGE-COPY	REBUILD INDEX	REORG INDEX	REORG TABLE-SPACE	UNLOAD
SYSPUNCH	Ignored	Ignored	Ignored	Ignored	Ignored	Ignored	Ignored	Ignored	NEW CATLG CATLG	NEW CATLG CATLG
SYSCOPY	Ignored	Ignored	NEW CATLG CATLG	Ignored	NEW CATLG CATLG	NEW CATLG CATLG	Ignored	Ignored	NEW CATLG CATLG	Ignored
SYSCOPY2	Ignored	Ignored	NEW CATLG CATLG	Ignored	NEW CATLG CATLG	NEW CATLG CATLG	Ignored	Ignored	NEW CATLG CATLG	Ignored
SYSRCPY1	Ignored	Ignored	NEW CATLG CATLG	Ignored	NEW CATLG CATLG	NEW CATLG CATLG	Ignored	Ignored	NEW CATLG CATLG	Ignored
SYSRCPY2	Ignored	Ignored	NEW CATLG CATLG	Ignored	NEW CATLG CATLG	NEW CATLG CATLG	Ignored	Ignored	NEW CATLG CATLG	Ignored
SYSUT1	NEW DELETE CATLG	NEW DELETE CATLG	Ignored	Ignored	NEW DELETE CATLG	Ignored	NEW DELETE CATLG	NEW CATLG CATLG	NEW DELETE CATLG	Ignored
SORTOUT	NEW DELETE CATLG	Ignored	Ignored	Ignored	NEW DELETE CATLG	Ignored	Ignored	NEW DELETE CATLG	NEW DELETE CATLG	Ignored
SYSMAP	Ignored	Ignored	Ignored	Ignored	NEW CATLG CATLG	Ignored	Ignored	Ignored	Ignored	Ignored
SYSERR	NEW CATLG CATLG	Ignored	Ignored	Ignored	NEW CATLG CATLG	Ignored	Ignored	Ignored	Ignored	Ignored

Table 98. Data dispositions for dynamically allocated data sets on RESTART

ddname	CHECK DATA	CHECK INDEX or CHECK LOB	COPY	COPY-TOCOPY	LOAD	MERGE-COPY	REBUILD INDEX	REORG INDEX	REORG TABLE-SPACE	UNLOAD
SYSREC	Ignored	Ignored	Ignored	Ignored	OLD KEEP KEEP	Ignored	Ignored	Ignored	MOD CATLG CATLG	MOD CATLG CATLG
SYSDISC	Ignored	Ignored	Ignored	Ignored	MOD CATLG CATLG	Ignored	Ignored	Ignored	MOD CATLG CATLG	Ignored
SYSPUNCH	Ignored	Ignored	Ignored	Ignored	Ignored	Ignored	Ignored	Ignored	MOD CATLG CATLG	MOD CATLG CATLG
SYSCOPY	Ignored	Ignored	MOD CATLG CATLG	Ignored	MOD CATLG CATLG	MOD CATLG CATLG	Ignored	Ignored	MOD CATLG CATLG	Ignored
SYSCOPY2	Ignored	Ignored	MOD CATLG CATLG	Ignored	MOD CATLG CATLG	MOD CATLG CATLG	Ignored	Ignored	MOD CATLG CATLG	Ignored
SYSRCPY1	Ignored	Ignored	MOD CATLG CATLG	Ignored	MOD CATLG CATLG	MOD CATLG CATLG	Ignored	Ignored	MOD CATLG CATLG	Ignored

Table 98. Data dispositions for dynamically allocated data sets on RESTART (continued)

ddname	CHECK DATA	CHECK INDEX or CHECK LOB	COPY	COPY-TOCOPY	LOAD	MERGE-COPY	REBUILD INDEX	REORG INDEX	REORG TABLE-SPACE	UNLOAD
SYSRCPY2	Ignored	Ignored	MOD CATLG CATLG	Ignored	MOD CATLG CATLG	MOD CATLG CATLG	Ignored	Ignored	MOD CATLG CATLG	Ignored
SYSUT1	MOD DELETE CATLG	MOD DELETE CATLG	Ignored	Ignored	MOD DELETE CATLG	Ignored	MOD DELETE CATLG	MOD CATLG CATLG	MOD DELETE CATLG	Ignored
SORTOUT	MOD DELETE CATLG	Ignored	Ignored	Ignored	MOD DELETE CATLG	Ignored	Ignored	MOD DELETE CATLG	MOD DELETE CATLG	Ignored
SYSMAP	Ignored	Ignored	Ignored	Ignored	MOD CATLG CATLG	Ignored	Ignored	Ignored	Ignored	Ignored
SYSERR	MOD CATLG CATLG	Ignored	Ignored	Ignored	MOD CATLG CATLG	Ignored	Ignored	Ignored	Ignored	Ignored

MODELDCB *dsname*
Specifies the DSN of a model data set. DCB information is read from the model data set.

BUFNO *integer*
Specifies the number of BSAM buffers. The specified value must be between 0 and 99. The **default** is **30**.

DATACLAS *name*
Specifies the SMS data class. The *name* value must be a valid SMS data class and must not exceed eight characters in length.

The data set will be cataloged if DATACLAS is specified. If this option is omitted, no DATACLAS is specified to SMS.

MGMTCLAS *name*
Specifies the SMS management class. The *name* value must be a valid SMS management class and must not exceed eight characters in length.

The data set will be cataloged if MGMTCLAS is specified. If this option is omitted, no MGMTCLAS is specified to SMS.

STORCLAS *name*
Specifies the SMS storage class. The *name* value must be a valid SMS storage class and must not exceed eight characters in length.

The data set is cataloged if STORCLAS is specified. If this option is omitted, no STORCLAS is specified to SMS.

RETPD *integer*
Specifies the retention period in days for the data set. The integer value must be between 1 to 9999.

If DATACLAS, MGMTCLAS or STORCLAS are specified, retention might be controlled by the class definition. RETPD cannot be specified with EXPDL and it is not valid with the **JES3DD** option.

EXPDL *'date'* Specifies the expiration date for the data set in the form YYYYDDD,

where YYYY is the four digit year and DDD is the three digit Julian day. The 'date' value must be enclosed by single quotes.

If DATACLAS, MGMTCLAS or STORCLAS are specified, retention may be controlled by the class definition. EXPDL may not be specified with RETPD and it is not valid with the JES3DD option.

VOLUMES *(vol1,vol2,...)*
Specifies a list of volume serial-numbers for this allocation. If the data set is not cataloged, the list will be truncated, if necessary, when stored in SYSIBM.SYSCOPY. The number of volumes specified may not exceed the specified or default value of VOLCNT.

The first volume must contain enough space for the primary space allocation.

VOLCNT *(integer)*
Specifies the maximum number of volumes that an output data set might require. The **default** value is **5** for DASD templates and **99** for TAPE templates.

GDGLIMIT *(integer)*
Specifies the number of entries to be created in a GDG base if a GDG DSN is specified and the base does not already exist. Specify a GDGLIMIT of zero (0) to not define a GDG base if one does not exist.

The **default** value is **99**. The integer value must be between 1 to 255.

UNCNT *integer*
Specifies the number of devices to be allocated. The specified value must be between 0 and 59.

If UNIT specifies a specific device number, then the value of UNCNT must be either 1 or omitted.

End of common-options

disk-options

SPACE *(primary,secondary)*
Specifies the MVS disk space allocation parameters between 1 to 1677215. If you specify *(primary,secondary)* value, these values are used instead of the DB2 calculated values. When specifying primary and secondary quantities, both values must be either specified or omitted.

Use the MAXPRIME option to set an upper limit on the *primary* quantity.

CYL
Specifies that allocation quantities, if present, are expressed in CYLINDERS and that allocation is to occur in cylinders. If SPACE CYL is specified, without (primary,secondary), the DB2 calculated quantities will be allocated in CYLINDERS using 3390 quantities for byte conversion.

The **default** is **CYL**.

TRK
If SPACE TRK is specified, without (primary,secondary), the

DB2 calculated quantities are allocated in tracks by using 3390 disk drive quantities for byte conversion.

MB
Specifies that allocation quantities, if present, are expressed in megabytes, and that allocation occurs in records. One megabyte is 1 048 576 bytes. If SPACE MB is specified, without (primary,secondary), the DB2-calculated quantities are allocated in records by using the average record length for the data set.

See "Default space calculations" on page 474 for default space calculations for each utility data set.

PCTPRIME *integer*
Specifies the percentage of the required space to be obtained as the primary quantity. The **default** is **100**.

Use the MAXPRIME option to set the upper limit of this value for large objects.

MAXPRIME *integer*
Specifies the maximum allowable primary space allocation, expressed in cylinders (CYL). This value constrains the *primary* space value and the PCTPRIME calculation.

NBRSECND *integer*
Specifies the division of secondary space allocations. After the primary space is allocated, the remaining space is divided into the specified number of secondary allocations. The integer value must be between 1 and 10. The **default** value is **10**.

End of disk-options

tape-options

STACK Specifies whether output data sets must be stacked contiguously on the same tape volumes.

NO
Do not stack output data sets contiguously on tape.

The **default** is **NO**.

YES
Stack similar output data sets as successive files on one logical tape volume. Within one utility execution, output data sets will be are stacked on a logical tape volume of the same usage type. For example, local primary image copies are stacked separately from local backup image copies.

To preserve parallel processing, parallel tasks are written to different tape volumes. The specific volume to which the data set is written can vary depending on the number of output data sets being produced, the number of parallel processes requested, and the number of tape units available to the job step.

An end-of-volume message terminates the stack, however the file that is currently being written when the EOV is received continues onto a second physical volume.

JES3DD *ddname*
Specifies the JCL DD name to be used at job initialization time for the tape unit. JES3 requires that all required tape units be pre-allocated by DD statement. Use the JES3DD to tell utility processing which unit to use for this template.

TRTCH Specifies the track recording technique for magnetic tape drives with improved data recording capability.

NONE
Specify TRTCH NONE to eliminate the TRTCH specification from dynamic allocation.

The **default** is **NONE**.

COMP
Specify TRTCH COMP to write data in compacted format.

NOCOMP
Specify TRTCH NOCOMP to write data in standard format.

End of tape-options

Instructions for using TEMPLATE

Some DB2 utilities produce data sets as a by-product or as an end result of utility execution. These data sets are referenced on utility control statements by a set of DD name keywords and are specified in detail on the corresponding JCL DD cards.

You can use the TEMPLATE utility control statement in place of DD cards to dynamically allocate utility data sets. With this utility control statement, you can establish data set policies, refer to those policies on utility control statements, and let DB2 utilities administer those policies at execution time. Templates contain:

- The data set naming convention
- DF/SMS parameters
- DASD or TAPE allocation parameters

You can specify a template in the SYSIN data set, immediately preceding the utility control statement that references it, or in one or more TEMPLATE data sets. The TEMPLATE data set DD name is specified on the OPTIONS utility control statement by TEMPLATEDD(*ddname*) and applies to all subsequent utility control statements until the end of input or until DB2 encounters a new OPTIONS TEMPLATEDD(*ddname*) specification. The **default** TEMPLATE data set DD name is **SYSTEMPL**.

TEMPLATE data sets can contain only TEMPLATE utility control statements. Any template that is defined within SYSIN overrides another template definition of the same name in a TEMPLATE data set.

With TEMPLATE utility control statement, database administrators can standardize data set allocation and the utility control statements that refer to those data sets, which reduces the need to customize and alter utility job streams.

You can create a library of TEMPLATE control statements. The default DD name of this library data set is SYSTEMPL, but you can change it by using the OPTIONS TEMPLATEDD control statement. This data set can contain only LISTDEF control

statements. If a TEMPLATE of the same name is found in SYSIN and in SYSTEMPL data sets, the name from SYSIN is used.

Key TEMPLATE operations

Like both LISTDEF and OPTIONS utility control statements, a TEMPLATE control statement performs a setup operation in preparation for use by another utility. When the control statement is ″processed″ the information is saved under the template name for the duration of the job-step. You can referenced it as though it were an output data set DD name by substituting the template name in place of the DD name on most utility control statements. If a DD name and a TEMPLATE name conflict, the DD card is used for allocation and the TEMPLATE is ignored. Minimally, a TEMPLATE statement consists of a name (similar to a DD name) and a data set naming convention. If no more is specified, DB2 calculates the required data set size and uses default data set attributes appropriate to the data set being created. DB2 then allocates a disk data set with these defaults.

The TEMPLATE statement required might look something like this:

```
TEMPLATE tmp1  DSN(DB2.&TS..D&JDATE..COPY&ICTYPE.&LOCREM.&PRIBAC.)
    VOLUMES(vol1,vol2,vol3)
LISTDEF payroll INCLUDE TABLESPACE PAYROLL.*
                INCLUDE INDEXSPACE PAYROLL.*IX
                EXCLUDE TABLESPACE PAYROLL.TEMP*
                EXCLUDE INDEXSPACE PAYROLL.TMPIX*
COPY LIST payroll  COPYDDN(tmp1,tmp1) RECOVERYDDN(tmp1,tmp1)
```

See “Syntax and options of the control statement” on page 463 for details.

Creating data set names

The data set naming convention that is specified on each TEMPLATE statement must be appropriate for the data set being created and coordinated with the other templates and DD cards in the same job-step. The data set name must be both unique and meaningful. Follow these guidelines when developing template names:

- Use a combination of static characters, national characters, and the provided variable names to form valid MVS data set qualifiers. Normal MVS rules apply. Variables that produce numeric values must be preceded by either a static character or a character variable. All qualifiers must start with an alphabetical character, and consist of a maximum of eight characters per qualifier and a maximum of 44 characters for the entire data set name.
- Use the two-character form of the DSN variables to save space.
- Use a double period ″..″ following all variables that precede the last qualifier (one to terminate the variable followed by a second static period to separate the qualifiers). For example: &DB..&TS.
- Use &DB. and &TS. to relate the data set to a data base object.
- Use P&PART. when executing PARTLEVEL lists.
- Use &JO. and &ST. to eliminate conflicts with other jobs or job steps.
- Use &SS. &US. &UT. and &UN. if it is critical to know the subsystem, member, user, utility-id, or utility name that produced the data set.
- Use D&DATE. and T&TIME. or the shorter substring variations to guarantee uniqueness.
- Use &IC. &LR. and &PB. to identify image copy data sets. For example, COPY&IC.&LR.&PB. would make a meaningful seven-character data set qualifier.

Controlling data set size

TEMPLATE syntax also allows user-specified disk space parameters. If you do not specify the SPACE keyword, DB2 estimates the size of the data set based on formulas that vary according to the utility and the data set. See “Default space calculations” for details.

In most cases, DB2 estimates the size of a data set based on the size of other existing data sets. Size can not be estimated if any of the required data sets are on tape. When size can be calculated, it is often a maximum size and can result in overly large data sets. Data sets are always allocated with the RLSE (release) option so that unused space is released on de-allocation. However in some cases, the calculated size of required data sets is too large for the DYNALLOC interface to handle. In this case error message DSNU1034I is issued, and the data set must be allocated by a DD card. If the object is part of a LISTDEF list, you might need to remove it from the list and process it individually.

Database administrators can check utility control statements without execution by using the PREVIEW function. In PREVIEW mode, DB2 expands all TEMPLATE data set names in the SYSIN DD in addition to any data set name from the TEMPLATE DD that are referenced on a utility control statement. DB2 then prints the information to the SYSPRINT data set and halts execution. You can specify PREVIEW in one of two ways, either as a JCL PARM or on the OPTIONS PREVIEW utility control statement.

Default space calculations

Three keywords are provided to let you manage how DB2 allocates the space required for the data set. You can use these three keywords, in combination, to constrain and quantify the allocation extents.

PCTPRIME
100% of the required space is allocated as a PRIMARY quantity. If this amount of space is typically not available on a single volume, decrease PCTPRIME. To avoid secondary extents keep PCTPRIME at 100.

MAXPRIME
If you want an upper limit based on size, and not percentage, use MAXPRIME.

NBRSECND
After the restrictions on the PRIMARY quantity have been applied, the remaining SECONDARY quantity is divided into the specified number of secondary extents.

If you omit the SPACE option quantities, current data set space estimation formulas shown in the “Data sets used by...” section of each online utility in this book are implemented as default values for DASD data sets.

Working with TAPE

The STACK keyword supports tape processing in two forms. The first form, STACK NO, supports traditional, single-file processing. The data set is written and the tape is rewound and repositioned or even remounted. The second form, STACK YES, lets successive files be written on a single logical tape without repositioning or remounting. Important considerations for STACK YES processing include:

- Only ″like files″ can be stacked on the same tape. For example, one tape might contain local primary image copies while another tape might contain remote primary image copies. The file types cannot be mixed.

- Files are only stacked within a single utility invocation. When that utility ends, the stack is terminated (the tape is rewound and unloaded). To allow stacking, use a LISTDEF list to force multiple objects to be processed under a single utility invocation.
- Parallel processing can complicate stacking. To prevent conflicts between parallel processes, only a single process can write a file to a given stack.
- Stacks are terminated at the physical end of volume (EOV). When EOV is reached, a second physical tape volume is called, and the data set in progress is completed; the tape is then rewound and unloaded. As a result, a logical tape volume can be contained on a maximum of two physical volumes.

Working with GDGs

TEMPLATE DSN operands support both GDG absolute version references and relative references. DB2 detects the absence of a GDG base and creates it, with a limit of 99 entries, by default. Use the keyword GDGLIMIT to alter this value or prohibit this action. If you use the PREVIEW function on the OPTIONS utility control statement, DB2 displays the GDG relative version references.GDG names are restricted to 35 characters.

A model data set as defined in the MODELDCB option might be required to allocate GDG data sets in your environment.

Concurrency and compatibility

TEMPLATE is a control statement that is used to setup an environment for another utility to follow. The template is stored until it is referenced by a specific utility. The list is expanded when it is referenced by another utility. At that time, the concurrency and compatibility restrictions of that utility apply, and the catalog tables necessary to expand the list must be available for read-only access.

Sample control statements

Example 1: A minimal template for an image copy, on disk, using default options.

```
TEMPLATE COPYDS DSN &DB..&TS..COPY&IC.&LR.&PB..D&DATE..T&TIME.
```

Example 2: Specify a tape template with an expiration date.

```
TEMPLATE TAPEDS DSN(&DB..&TS..D&DATE.)
    UNIT 3590-1  EXPDL '2100001'
```

Example 3: Specify a disk template with SPACE.

```
TEMPLATE DISK DSN &DB..&TS..T&TIME.
    SPACE(100,10) CYL
```

Example 4: Specify a disk template using a default size with constraints.

```
TEMPLATE DISK DSN(&DB..&TS..T&TIME.)
    SPACE CYL  MAXPRIME 1000
```

Example 5: Use template TMP1 to define the naming convention for the COPYDDN data set.

```
LISTDEF CPY1 INCLUDE TABLESPACES TABLESPACE DBA906*.T*A906*
             INCLUDE INDEXSPACES COPY YES INDEXSPACE ADMF001.I?A906*
TEMPLATE TMP1 UNIT SYSDA
```

```
                    DSN (DH109006.&STEPNAME..&SN..T&TIME.)
                    DISP (MOD,CATLG,CATLG)
COPY LIST CPY1 COPYDDN (TMP1) PARALLEL (2) SHRLEVEL REFERENCE
```

Note: Parentheses for the DSN name-expression are optional.

Example 6: Use TEMPLATE to create a GDG data set.

```
//*************************************************************
//* COMMENT:   Define a model dataset.                        *
//*************************************************************
//STEP1    EXEC   PGM=IEFBR14
//SYSCOPX  DD DSN=JULTU225.MODEL,DISP=(NEW,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20)),VOL=SER=SCR03,
//         DCB=(RECFM=FB,BLKSIZE=4000,LRECL=100)
//SYSPRINT DD  SYSOUT=A
//SYSIN    DD  *
/*
//*************************************************************
//* COMMENT: GDGLIMIT(6)
//*************************************************************
//STEP2    EXEC DSNUPROC,UID='JULTU225.GDG',
//         UTPROC='',
//         SYSTEM='SSTR'
//SYSUT1   DD DSN=JULTU225.GDG.STEP2.SYSUT1,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SORTOUT  DD DSN=JULTU225.GDG.STEP2.SORTOUT,
//         DISP=(MOD,DELETE,CATLG),
//         UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSIN    DD *
   TEMPLATE COPYTEMP
            UNIT SYSDA
            DSN 'JULTU225.GDG(+1)'
            MODELDCB JULTU225.MODEL
            GDGLIMIT(6)
   COPY TABLESPACE DBLT2501.TPLT2501
        FULL YES
        COPYDDN (COPYTEMP)
        SHRLEVEL REFERENCE
/*
```

Chapter 29. UNLOAD

The UNLOAD online utility unloads data from one or more source objects to one or more BSAM sequential data sets in external formats. The source can be DB2 table spaces or DB2 image copy data sets. UNLOAD must be run on the system where the definitions of the table space and the tables exist.

UNLOAD is an enhancement of REORG UNLOAD EXTERNAL. With UNLOAD, you can unload rows from an entire table space or select specific partitions or tables to unload. You can also select columns by using the field specification list. If a table space is partitioned, you can unload all of the selected partitions into a single data set, or you can unload each partition in parallel into physically distinct data sets.

The output records written by the UNLOAD utility are compatible as input to the LOAD utility, as a result, you can reload the original table or different tables.

For a diagram of UNLOAD syntax and a description of available options, see "Syntax and options of the control statement". For detailed guidance on running this utility, see "Instructions for running UNLOAD" on page 507.

Output: Output from UNLOAD consists of one or more of the following:
- An unloaded table space or partition
- A discard file of rejected records

Authorization required: To execute this utility, the privilege set of the process must include one of the following:
- Ownership of the tables
- SELECT privilege on the tables
- DBADM authority for the database
- SYSADM authority
- SYSCTRL authority (catalog tables only)

Execution phases of UNLOAD: The UNLOAD utility operates in these phases:

Phase	Description
UTILINIT	Initialization and setup.
UNLOAD	Unloading records to sequential data sets. One pass through the input data set is made. If UNLOAD is processing a table space or partition, DB2 takes internal commits to provide commit points at which to restart in case operation should halt in this phase.
UTILTERM	Cleanup.

Syntax and options of the control statement

The utility control statement defines the function the utility job performs. You can create a control statement with the ISPF/PDF edit function. After creating it, save it in a sequential or partitioned data set. When you create the JCL for running the job, use the SYSIN DD statement to specify the name of the data set that contains the utility control statement.

Syntax diagram

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

© Copyright IBM Corp. 1983, 2001

UNLOAD

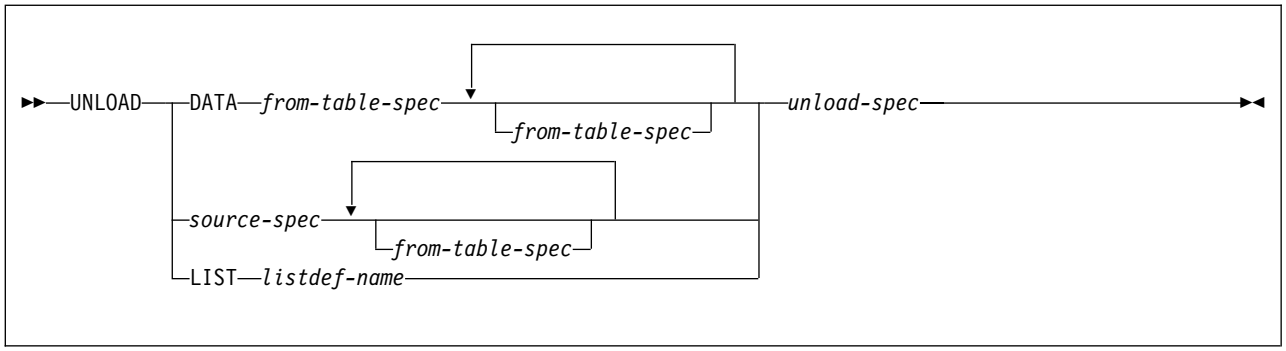

source-spec:

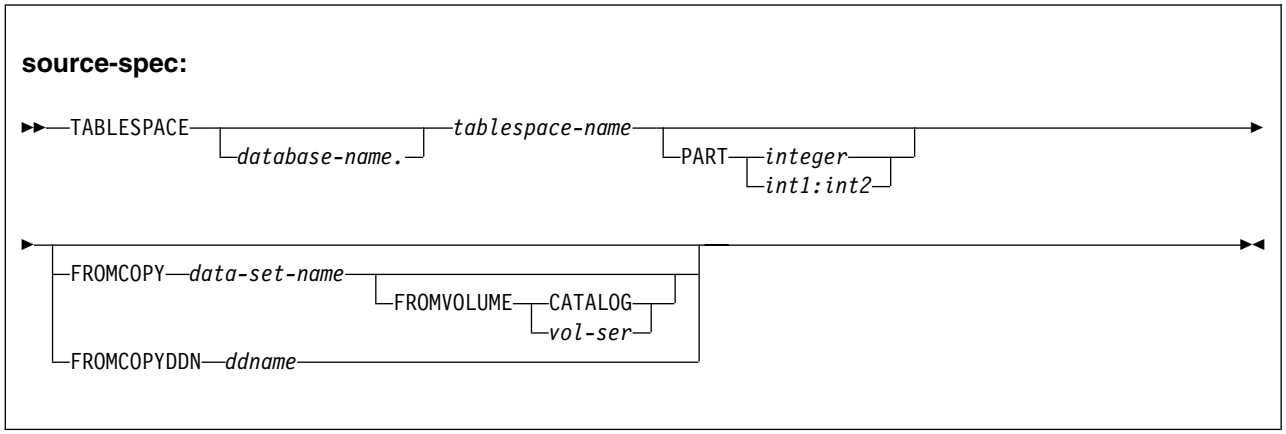

unload-spec:

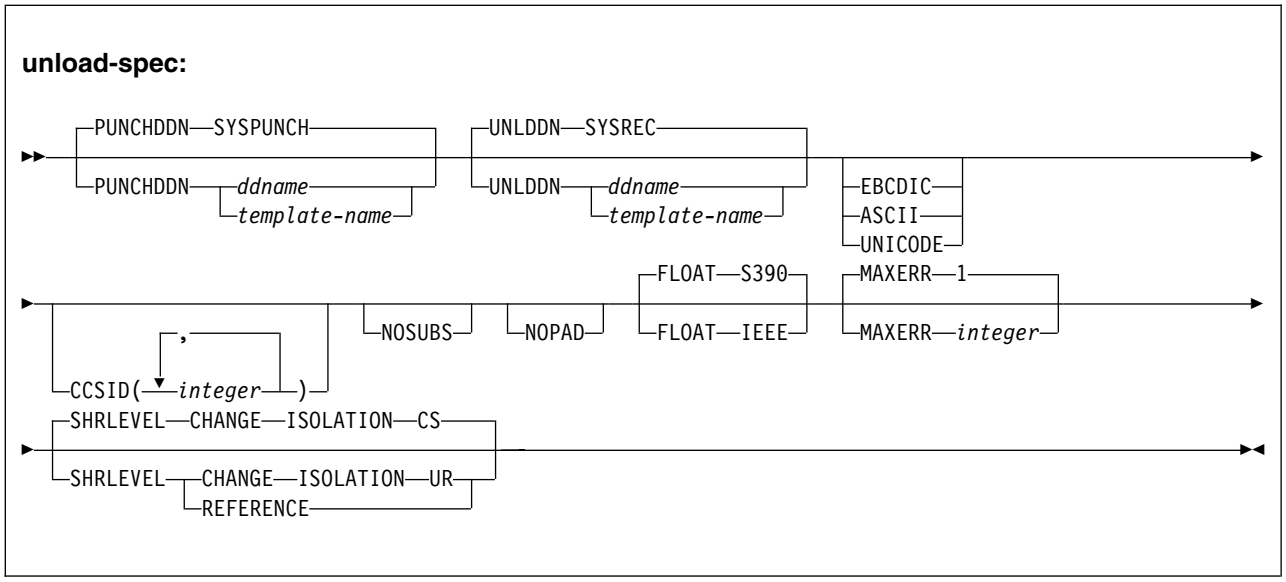

For the syntax diagram and the option descriptions of the FROM TABLE specification, see “FROM-TABLE-spec” on page 485.

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

DATA Identifies the data selected for unloading with *table-name* in the

from-table-spec. The DATA keyword is mutually exclusive with TABLESPACE, PART and LIST keywords and *from-copy-spec*.

When you specify the DATA keyword, or you omit either the TABLESPACE or the LIST keyword, you must also specify at least one FROM TABLE clause.

TABLESPACE
Specifies the table space (and, optionally, the database to which it belongs) from which the data is unloaded.

database-name
The name of the database to which the table space belongs. The name can not be DSNDB01 or DSNDB07.

The **default** is **DSNDB04**.

tablespace-name
The name of the table space from which the data is unloaded. The specified table space must not be a LOB table space.

PART Identifies a partition or a range of partitions from which the data is unloaded. This keyword applies when the specified table space is partitioned; it cannot be specified with LIST. The maximum is 254.

integer
Designates a single partition. *integer* must be an existing partition number within the table space.

int1:int2
Designates a range of partitions from *int1* to *int2*. *int1* must be a positive integer that is less than the highest partition number within the table space. *int2* must be an integer greater than *int1* and less than or equal to the highest partition number.

If no PART keyword is specified in an UNLOAD control statement, the data from the entire table space will be unloaded into a single unload data set.

FROMCOPY *data-set-name*
Indicates that data is unloaded from an image copy data set. When you specify FROMCOPY, the UNLOAD utility processes only the specified image copy data set. Alternatively, the FROMCOPYDDN keyword can be used where multiple image copy data sets can be concatenated under a single DD name.

dataset-name
Specifies the name of a single image copy data set.

The FROMCOPY image copy data set must have been created by one of the following utilities:
- COPY
- LOAD inline image copy
- MERGECOPY
- REORG TABLESPACE inline image copy
- DSN1COPY

If the specified image copy data set is a full image copy, either compressed or uncompressed records can be unloaded.

If the specified image copy data set is an incremental image copy or a copy of a partition or partitions, you can unload compressed records

only when the same data set contains the dictionary pages for decompression. If an image copy data set contains a compressed row and a dictionary is not available, then DB2 issues an error message. See the MAXERR option on page 484 for more information about specifying error tolerance conditions.

When you specify FROMCOPY or FROMCOPYDDN, you can also specify selection criteria with either PART, or FROM TABLE, or both, to qualify tables and rows to be unloaded.

FROMVOLUME
Identifies the image copy data set.

CATALOG
Identifies that the data set is cataloged. Use this option only for an image copy that was created as a cataloged data set (its volume serial is not recorded in SYSIBM.SYSCOPY).

vol-ser
Identifies the data set by an alphanumeric volume serial identifier of its first volume. Use this option only for an image copy that was created as a non-cataloged data set. To locate a data set stored on multiple tape volumes, specify the first *vol-ser* in the SYSCOPY record .

FROMCOPYDDN *ddname*
Indicates that data is unloaded from one or more image copy data sets associated with the specified DDNAME. Multiple image copy data sets (primarily for copy of pieces) can be concatenated under a single DD name.

ddname
Identifies a DD name with which one or more image copy data sets are associated.

LIST *listdef-name*
Identifies the name of a list of objects that are defined by a LISTDEF utility control statement. Index spaces, LOB table spaces, and directory objects must not be included in the list. You cannot use the LIST option to specify image copy data sets.

When you specify the LIST option, the referenced LISTDEF identifies:
- The table spaces from which the data is unloaded (you can use the pattern-matching feature of LISTDEF).
- The partitions (if a table space is partitioned) from which the data is unloaded (defined by the INCLUDE, EXCLUDE and PARTLEVEL keywords in the LISTDEF statement).

The UNLOAD utility associates a single table space with one output data set, except when partition-parallelism is activated. When you use the LIST option with a LISTDEF that represents multiple table spaces, you must also define a data set TEMPLATE that corresponds to all the table spaces and specify the *template-name* in the UNLDDN option.

If you want to generate the LOAD statements, you must define another TEMPLATE for the PUNCHDDN data set that is similar to UNLDDN. DB2 then generate a LOAD statement for each table space.

PUNCHDDN

Specifies the DD name for a data set, or a template name, that defines one or more data set names to receive the LOAD utility control statements that the UNLOAD utility generates.

ddname
Specifies the DD name.

The **default** is **SYSPUNCH**.

template-name
Identifies the name of a data set template that is defined by a TEMPLATE utility control statement.

If the name is defined as a DD name by the JCL, it is treated as the DD name.

When you run the UNLOAD utility for multiple table spaces and you want to generate corresponding LOAD statements, you must have multiple output data sets that correspond to the table spaces so that DB2 retains all of the generated LOAD statements. In this case, you must specify an appropriate template name to PUNCHDDN. If you omit the PUNCHDDN specification, the LOAD statements are not generated.

If the partition variable (&PART. or &PA.) is included in a TEMPLATE for PUNCHDDN, DB2 replaces the &PART. or &PA. variable with the lowest partition number in the list of partitions to be unloaded. The partition number is in the form *nnnnn*.

UNLDDN

Specifies the DD name for a data set or a template name that defines one or more data set names into which the data is unloaded.

ddname
Specifies the DD name.

The **default** is **SYSREC**.

template-name
Identifies the name of a data set template that is defined by a TEMPLATE utility control statement.

If the specified name is defined both as a DDNAME (in the JCL) and as a template name (a TEMPLATE statement), it is treated as the DDNAME.

When you run the UNLOAD utility for a partitioned table space, the selected partitions are unloaded in parallel if a template name is specified in UNLDDN and the template data set name contains the partition as a variable (&PART. or &PA.). In this case, the template is expanded into multiple data sets that correspond to the selected partitions. For example, the data from partitions 2 and 5 of the table space *testdb.testtsp1* are unloaded in parallel into two data sets by using the following utility statement:

```
TEMPLATE unldtempl
    DSNAME(&DBNAME..&TSNAME..P&PART.)
LISTDEF unldlist
    INCLUDE TABLESPACE testdb.testtsp1 PARTLEVEL (2)
    INCLUDE TABLESPACE testdb.testtsp1 PARTLEVEL (5)
UNLOAD LIST unldlist UNLDDN unldtempl ...
```

Similarly, when you run the UNLOAD utility for multiple table spaces, the output records are placed in data sets that correspond to the respective table spaces; therefore the output data sets must be physically distinctive, and you must specify an appropriate template name to UNLDDN. If you omit the UNLDDN specification, the SYSREC DDNAME is not used and an error occurs.

If the partition variable (&PART. or &PA.) is included in TEMPLATE DSNAME when the partition parallelism is not applicable (when the source is a non-partitioned table space or an image copy), the variable is replaced by '00000' in the actual data set name. In this case, warning message DSNU1252I is issued, and the UNLOAD utility issues return code 4.

EBCDIC
Specifies that all output data of the character type will be in EBCDIC. If a different encoding scheme is used for the source data, the data is converted into EBCDIC (except for bit strings).

If you do not specify either EBCDIC, ASCII, UNICODE, or CCSID, the encoding scheme of the source data is preserved.

See the description of the CCSID option for this utility.

ASCII Specifies that all output data of the character type will be in ASCII. If a different encoding scheme is used for the source data, the data is converted into ASCII (except for bit strings).

If you do not specify either EBCDIC, ASCII, UNICODE, or CCSID, the encoding scheme of the source data is preserved.

See the description of the CCSID option for this utility.

UNICODE
Specifies that all output data of the character type will be in UNICODE (except for bit strings). If a different encoding scheme is used for the source data, the data is converted into UNICODE.

If you do not specify either EBCDIC, ASCII, UNICODE, or CCSID, the encoding scheme of the source data is preserved.

See the description of the CCSID option of this utility.

CCSID *(integer1,integer2,integer3)*
Specifies up to three coded character set identifiers (CCSIDs) to be used for the data of character type in the output records, including data unloaded in the external character formats.

integer1 specifies the CCSID for SBCS data. *integer2* specifies the CCSID for mixed data. *integer3* specifies the CCSID for DBCS data. This option is not applied to data with a subtype of BIT.

The following specifications are also valid:

CCSID(*integer1***)**
Only an SBCS CCSID is specified.

CCSID(*integer1***,***integer2***)**
An SBCS CCSID and a mixed CCSID are specified.

integer
Valid CCSID specification or 0.

If you specify an argument as 0 or omit an argument, the encoding scheme specified by EBCDIC, ASCII, or UNICODE is assumed for the corresponding data type (SBCS, MIXED, or DBCS). If you do not specify either EBCDIC, ASCII, or UNICODE:

- If the source data is of character type, the original encoding scheme is preserved.
- For character strings that are converted from numeric data, date, time, or timestamp, the default encoding scheme of the table is used. For more information, see the CCSID option of the CREATE TABLE statement in Chapter 5 of *DB2 SQL Reference*.

When a CCSID conversion is requested, CCSID character substitutions can occur in the output string. Use the NOSUBS option to prevent possible character substitutions during CCSID conversion.

NOSUBS

Specifies that CCSID code substitution is not to be performed during unload processing.

When a string is converted from one CCSID to another (including EBCDIC, ASCII, and UNICODE), a substitution character is sometimes placed in the output string. For example, this substitution occurs when a character (referred to as a codepoint) that exists in the source CCSID does not exist in the target CCSID. You can use the NOSUBS keyword to prevent the UNLOAD utility from allowing this substitution.

If you specify the NOSUBS keyword and character substitution is attempted while unloading data, it is treated as a conversion error. The record with the error is not unloaded, and the process continues until the total error count reaches the number specified by MAXERR.

NOPAD

Specifies that the variable length columns in the unloaded records occupy the actual data length without additional padding. As a result, the unloaded or discarded records might have varying lengths.

When you do not specify NOPAD:

- Default UNLOAD processing pads variable length columns in the unloaded records to their maximum length, and the unloaded records have the same length for each table.
- The padded data fields are preceded by the length fields that indicate the size of the actual data without the padding.
- When the output records are reloaded using the LOAD utility, padded data fields are treated as varying length data.

While LOAD processes records with variable length columns that are unloaded or discarded by using the NOPAD option, these records can not be processed by applications that only process fields in fixed positions. For example, the LOAD statement generated for the EMP sample table would look similar to the LOAD statement shown in Figure 19 on page 348.

FLOAT

Specifies the output format of the numeric floating point data. This option applies to the binary output format only.

S390

Indicates that the binary floating point data is written to the output records in the S/390® internal format (also known as the “hexadecimal floating point” or HFP).

The **default** is **FLOAT S390**.

IEEE
Indicates that the binary floating point data is written to the output records in the IEEE format (also known as the "binary floating point" or BFP).

The IEEE option is applicable only when OS/390 Version 2 Release 6 or a subsequent version is installed, and a G5 or above processor is present.

MAXERR *integer*
Specifies the maximum number of records in error that are allowed; the unloading process terminates when this value is reached.

integer
Specifies the number of records in error that are allowed. When the error count reaches this number, the UNLOAD utility issues message DSNU1219 and terminates with return code 8.

The **default** is **1**, which indicates that UNLOAD stops when the first error is encountered. If you specify 0 or any negative number, execution continues regardless of the number of records in error.

If multiple table spaces are being processed, the number of records in error is counted for each table space. If the LIST option is used, you can add OPTION utility control statement (EVENT option with ITEMERROR) before the UNLOAD statement to specify that the table space in error is skipped and the subsequent table spaces are processed.

SHRLEVEL
Specifies whether other processes can access or update the table space or partitions while the data is being unloaded.

UNLOAD ignores the SHRLEVEL specification when the source object is an image copy data set.

The **default** is **SHRLEVEL CHANGE ISOLATION CS**.

CHANGE
Specifies that rows can be read, inserted, updated, and deleted from the table space or partition while the data is being unloaded.

ISOLATION
Specifies the isolation level with SHRLEVEL CHANGE.

CS
Indicates that the UNLOAD utility reads rows in cursor stability mode. With CS, the UNLOAD utility assumes CURRENTDATA(NO).

UR
Indicates that uncommitted rows, if they exist, are unloaded. The unload operation is performed with minimal interference from the other DB2 operations that are applied to the objects from which the data is being unloaded.

REFERENCE
Specifies that during the unload operation, rows of the tables can be read, but can not be inserted, updated, nor deleted by other DB2 threads.

When you specify SHRLEVEL REFERENCE, the UNLOAD utility drains writers on the table space from which the data is to be unloaded. When data is unloaded from multiple partitions, the drain lock will be obtained for all of the selected partitions in the UTILINIT phase.

FROM-TABLE-spec

More than one table or partition for each table space can be unloaded with a single invocation of the UNLOAD utility. One FROM TABLE statement for each table to be unloaded is required to identify:

- A table name from which the rows will be unloaded
- A field to identify the table that is associated with the rows that are unloaded from the table by using the HEADER option
- Sampling options for the table rows
- A list of field specifications for the table that is used to select columns to be unloaded
- Selection conditions, specified in the WHEN clause, that are used to qualify rows to be unloaded from the table

All tables that are specified by FROM TABLE statements must belong to the same table space. If rows from specific tables are to be unloaded, a FROM TABLE clause must be specified for each source table. If you do not specify a FROM TABLE clause for a table space, all the rows of the table space are unloaded.

A list of field specifications is used to specify the following characteristics:

- Column selection. Specifies the column names of a table to be unloaded. If a list of field specifications is given, only the listed columns will be unloaded.
- Column ordering. Specifies the order of fields placed in the output records. If a list of field specifications is given, data of the listed columns will be unloaded in the order of listed column names.
- Output field attributes and format. Specifies the data type, length and format of the data in the output records.

If you omit a list of field specifications, all columns of the source table are unloaded in the defined column order for the table. The default output field types corresponding to the data types of the columns are used.

In a FROM TABLE clause, you can use parenthesis in only two instances: to enclose the entire field selection list and in a WHEN selection clause. This usage avoids potential conflict between the keywords and field-names that are used in the field selection list. A valid sample of a FROM TABLE clause specification follows:

```
UNLOAD ...
  FROM TABLE tablename SAMPLE x (c1,c2) WHEN (c3>0)
```

You cannot specify FROM TABLE if the LIST option is already specified.

UNLOAD

FROM-TABLE-spec:

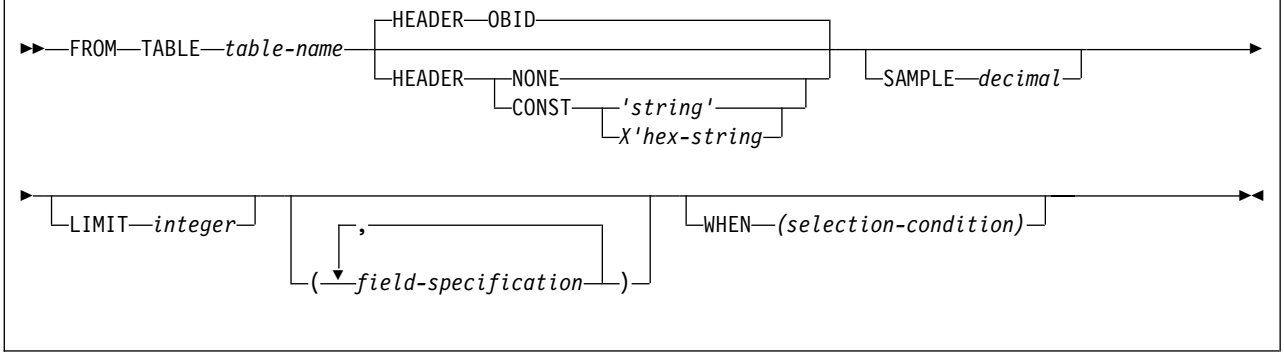

field-specification:

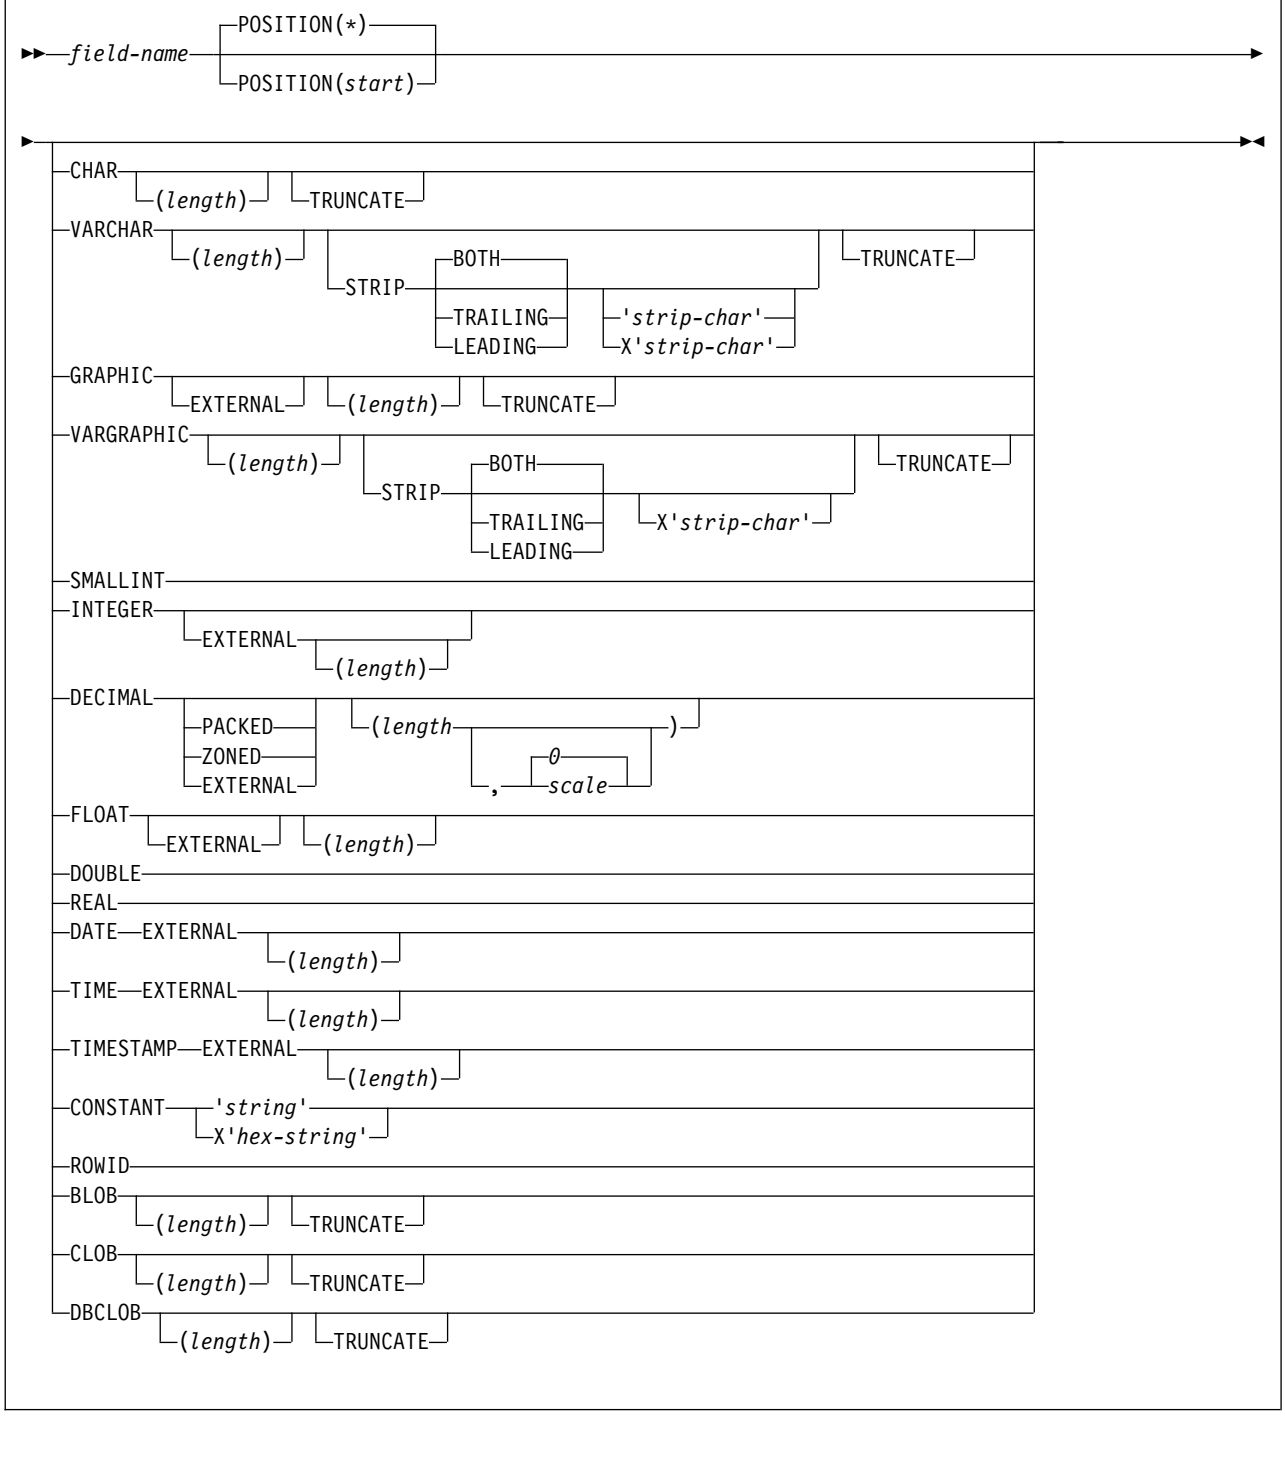

selection condition:

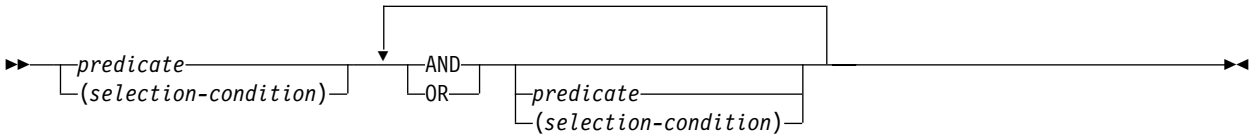

predicates:

basic predicate:

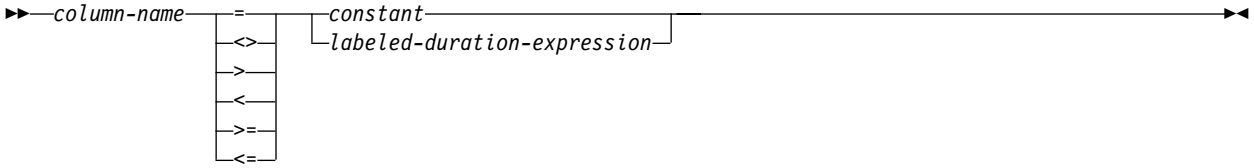

BETWEEN predicate:

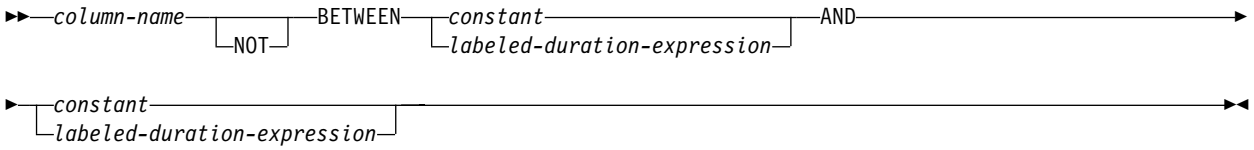

IN predicate:

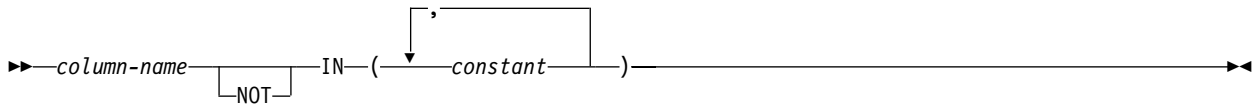

LIKE predicate:

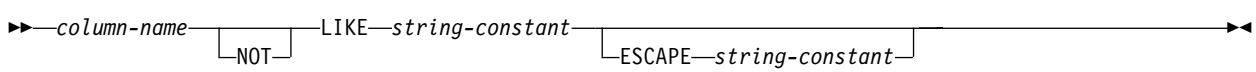

NULL predicate:

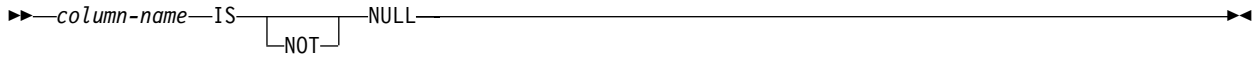

labeled-duration-expression:

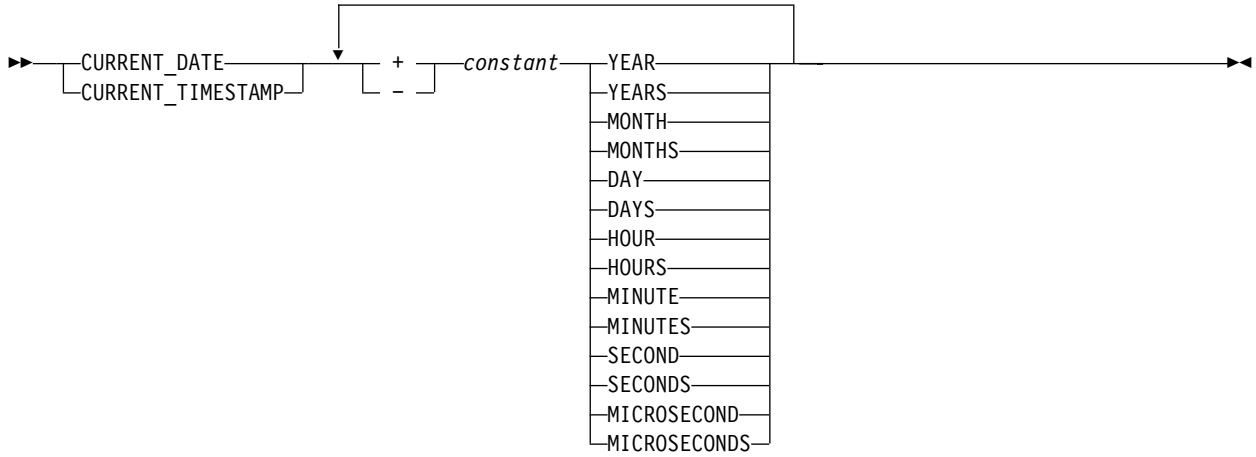

Option descriptions for FROM TABLE

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 23.

table-name

Identifies a DB2 table from which the rows are unloaded and to which the options in the FROM TABLE clause are applied.

If the table name is not qualified by an authorization ID, the authorization ID of the invoker of the utility job step is used as the qualifier of the table name.

HEADER

Specifies a constant header field, at the beginning of the output records, that can be used to associate an output record with the table from which it was unloaded.

If you specify a header field, it will be used as the field selection criterion of the WHEN clause (a part of the INTO-TABLE specification) in the LOAD statement that is generated.

OBID

Specifies that the object identifier (OBID) for the table (a two byte binary data) is placed in the first two bytes of the output records unloaded from the table.

If you omit the HEADER option, **HEADER OBID** is the default.

With HEADER OBID, the first two bytes of the output record cannot be used by the unloaded data. To illustrate, the following UNLOAD statement:

```
UNLOAD ...
  FROM TABLE tablename HEADER OBID ...
```

generates a LOAD statement similar to the following example:

```
LOAD ...
  INTO TABLE tablename WHEN (1:2)=X'hh' ...
```

where X'hh' is the hexadecimal notation of the OBID of table *tablename*.

NONE

Indicates that no record header field is created.

If HEADER NONE is specified in a FROM TABLE clause, the corresponding INTO TABLE clause in the generated LOAD statement will not have a WHEN specification. Therefore, if rows from multiple tables are unloaded and HEADER NONE is specified in one or more FROM TABLE clauses, rows unloaded from those tables will not be able to be reloaded until you edit the generated LOAD statement. If you use the generated statement directly with the LOAD utility, the results might be unpredictable.

CONST

Specifies that a constant string is used as the record header. The given string operand determines the length of the header field. The string value must be enclosed by a pair of single quote characters.

For example, the following UNLOAD statement

```
UNLOAD ...
  FROM TABLE tablename HEADER CONST 'abc' ...
```

generates a LOAD statement similar to the following example:

```
LOAD ...
  INTO TABLE tablename WHEN (1:3)='abc' ...
```

In this example, the given string is assumed to be in SBCS EBCDIC. The output string of the HEADER field is in the specified or the default

encoding scheme. If the encoding scheme used for output is not EBCDIC, the SBCS CCSID conversion is applied to the given string before it is placed in the output records. If the output SBCS encoding scheme is not EBCDIC, the WHEN condition in the generated LOAD statement contains a hex string.

You can also use the hexadecimal form, X'*hex-string*', to represent a string constant. If you want to specify a CONST string value in an encoding scheme other than SBCS EBCDIC, use the hexadecimal form. No CCSID conversion is performed if the hexadecimal form is used.

SAMPLE *decimal*

Indicates that only sampled rows of the table are unloaded. If selection conditions are specified by a WHEN clause within the same FROM TABLE clause, sampling is applied to the rows that are qualified by the WHEN selection conditions.

decimal

Specifies the percentage of the rows to be sampled in the decimal format. The precision is *ddd.dddd* and the valid range is *0 <= decimal <= 100*.

If the number of rows to which the sampling is applied is *N*:

- *decimal* × *N* / *100* rows are unloaded (the fraction might be rounded up to the nearest whole number).
- If *decimal > 0* and *N > 0* then at least one row is unloaded.
- If *decimal = 100*, all the rows from the table are unloaded.
- If the given *decimal = 0* or *N = 0*, no row will be unloaded from the table.

The sampling is applied per individual table. If the rows from multiple tables are unloaded with sampling enabled, the referential integrity between the tables might be lost.

LIMIT *integer*

Specifies the maximum number of rows to be unloaded from a table. If the number of unloaded rows reaches the specified limit, message DSNU1201 is issued for the table and no more rows are unloaded from the table. The process continues to unload qualified rows from the other tables.

When partition parallelism is activated, the LIMIT option is applied to each partition instead of the entire table.

integer

Indicates the maximum number of rows to be unloaded from a table. If the specified number is less than or equal to zero, no row is unloaded from the table.

Like the SAMPLE option, if multiple tables are unloaded with the LIMIT option, the referential integrity between the tables may be lost.

field-name

Identifies a column name that must exist in the source table.

POSITION *(start)*

Specifies the field position in the output record. The operand (position parameter) must be enclosed by parentheses.

The position parameter can be specified as:

***** An asterisk, indicating that the field starts at the first byte after the last position of the previous field.

start A positive integer that indicates the start column of the data field.

The **default** is **POSITION(*)**.

The POSITION option specification cannot be enclosed in parentheses; however, the *start* parameter must be enclosed in parentheses. For example

```
field-name POSITION(*)     -- a valid example
field-name POSITION(20)    -- a valid example
```

The first column (byte position) of an output record corresponds to POSITION(1). If you specify HEADER NONE in the FROM TABLE clause, the item specified by the HEADER option is placed at the beginning of all the records unloaded from the table. You must account for the space for the record header:

- HEADER OBID (the default case): two bytes from position 1.
- HEADER CONST 'string' or X'*hex-string*' case: The length of the given string from position 1.

If the source table column is nullable, a NULL indicator byte is prepended to the data field in the output record. The *start* parameter (or *) points to the position of the NULL indicator byte. In this situation the *start* is shifted by one byte to the right (as *start+1*) in the generated LOAD statement because in the LOAD statement, the start parameter of the POSITION option points to the next byte past the NULL indicator byte.

For a varying length field, a length field precedes the actual data field (after the NULL indicator byte, if nullable). If the field is not nullable, the *start* parameter (or *) points to the first byte of the length field. The size of the length field is either four bytes (BLOB, CLOB, or DBCLOB) or two bytes (VARCHAR or VARGRAPHIC).

When the output field positions are explicitly specified by using *start* parameters (or using the * format) of the POSITION option, the following items must be taken into account as a part of the output field:

- For a nullable field, a space for the NULL indicator byte.
- For varying length data, a space for the length field (either two bytes or four bytes).

"Determining the layout of output fields" on page 514 illustrates the field layout in conjunction with the POSITION option, NULL indicator byte, the length field for a varying length field, the *length* parameter, and the actual data length.

The POSITION option will be useful when the output fields must be placed at desired positions in the output records. The use of the POSITION parameters, however, can restrict the size of the output data fields. Use care when explicitly specifying *start* parameters for nullable and varying length fields. The TRUNCATE option might be required, if applicable, to fit a data item in a shorter space in an output record.

If you omit the POSITION option for the first field, the field will start from position 1 if HEADER NONE is specified. Otherwise, the field will start from the next byte position past the record header field. If POSITION is omitted for a subsequent field, the field is placed next to the last position of the previous field without any gap.

If NOPAD is specified and POSITION parameters are given for certain fields, the effect of the NOPAD option might be lost because the fields with *start* parameters (other than the default *) always start at the fixed positions in the output records.

CHAR Indicates that the output field is a character type with fixed length. You can use CHARACTER in place of CHAR. If the source table column is nullable, a NULL indicator byte is placed at the beginning of the output field.

If you specified the EBCDIC, ASCII, UNICODE, or CCSID options, the output data that corresponds to the specified option, is encoded in the CCSID, depending on the subtype of the source data (SBCS or MIXED). If the subtype is BIT, no conversion will be applied.

(length)
Specifies that size of the output data in bytes.

If the *length* parameter is omitted, the default is the maximum length defined on the source table column. When the *length* parameter is specified:

- If the *length* is less than the size of the table column, the data is truncated to the *length* if the TRUNCATE keyword is present; otherwise, a conversion error will occur.
- If the *length* is larger than the size of the table column, the output field is padded by the default pad characters to the specified length.

TRUNCATE
The optional TRUNCATE keyword indicates that a character string (encoded for output) will be truncated from the right, if the data does not fit in the available space for the field in the output record. Truncation occurs at the character boundary. See “Specifying TRUNCATE and STRIP options for output data” on page 516 for the truncation rules used in the UNLOAD utility. Without TRUNCATE, an error will occur when the output field size is too small for the data.

VARCHAR
Specifies that the output field type is character of varying length. A two-byte binary field indicating the length of data in bytes is prepended to the data field. If the table column is nullable, a NULL indicator byte is placed before the length field.

If you specified the EBCDIC, ASCII, UNICODE, or CCSID options, the output data is encoded in the CCSID corresponding to the specified option, depending on the subtype of the source data (SBCS or MIXED). Note that, if the subtype is BIT, no conversion will be applied.

(length)
Specifies the maximum length of the actual data field in bytes. If you also specify NOPAD, it indicates the maximum allowable space for the data in the output records; otherwise, the space of the specified length is reserved for the data.

If the length parameter is omitted, the default is the smaller of 255 and the maximum length defined on the source table column.

STRIP
With this option, blanks (the default) or the specified characters are removed from the beginning, the end, or both ends of the data, and the VARCHAR length field is adjusted to the length of the stripped data.

The STRIP option is applicable if the subtype of the source data is BIT. In this case, no CCSID conversion is performed on the specified strip-char (even if it is given in the form *'strip-char'*).

The effect of the STRIP option is the same as the SQL STRIP scalar function. For details, see Chapter 5 of *DB2 SQL Reference*.

BOTH
Indicates that occurrences of blank or the specified character are removed from the both ends. The **default** is **BOTH**.

TRAILING
Indicates that occurrences of blank or the specified character are removed from the end.

LEADING
Indicates that occurrences of blank or the specified character are removed from the beginning.

X'strip-char'
Specifies a single byte character to be stripped. It can be specified in the hexadecimal form, X'*hex-string*', where *hex-string* is two hexadecimal characters representing an SBCS character. If the *strip-char* operand is omitted, the default is the blank character (for the EBCDIC-encoded output case, X'40', for the ASCII-encoded output case, X'20', and for the UNICODE-encoded output case, X'20').

The strip operation is applied after the character code conversion, if the output character encoding scheme is different from the one defined on the source data. Therefore, if a strip character is specified in the hexadecimal format, it must be in the encoding scheme used for output.

'strip-char'
Specifies a single byte character to be stripped. A value given in the form *'strip-char'* is assumed to be in EBCDIC. Depending on the output encoding scheme, SBCS CCSID conversion is applied to the strip-char value before it is used in the strip operation. If you want to specify a strip-char value in an encoding scheme other than EBCDIC, use the hexadecimal form. No CCSID conversion is performed if the hexadecimal form is used.

TRUNCATE
The optional TRUNCATE keyword indicates that a character string (encoded for output) will be truncated from the right, if the data does not fit in the available space for the field in the output records. Truncation occurs at a character boundary. See “Specifying TRUNCATE and STRIP options for output data” on page 516 for the truncation rules used in the UNLOAD utility. Without TRUNCATE, an error will occur when the output field size is too small for the data.

GRAPHIC
Specifies that the output field is of the fixed length graphic type. If the table column is nullable, a NULL indicator byte is placed before the actual data field.

If the output is in EBCDIC, the shift-in and shift-out characters will not be included at the beginning and at the end of the data.

(length)
If you specify the *length* parameter, it must be the number of DBCS characters (the size of the output data in bytes is twice the given length). If the given *length* is larger than the source data length, the output field will be padded with the default pad character.

TRUNCATE
The optional TRUNCATE keyword indicates that a graphic character string (encoded for output) will be truncated from the right, if the data does not fit in the available space for the field in the output records. Truncation occurs at a character (DBCS) boundary. Without TRUNCATE, an error will occur when the output field size is too small for the data.

GRAPHIC EXTERNAL
Specifies that the data is written in the output records as a fixed length field of the graphic type with the external format; that is, the shift-out (SO) character is placed at the starting position, and the shift-in (SI) character at the ending position. The output field size will be even in the byte count.

GRAPHIC EXTERNAL is supported only in the EBCDIC output mode (by default or when the EBCDIC keyword is specified).

If the *start* parameter of the POSITION option is used to specify the output column position, it will point to the (inserted) shift-out character at the beginning of the field. The shift-in character is placed at the next byte position past the last double-byte character of the data.

(length)
If you specify the optional *length* parameter, it must be a number of DBCS characters not including the shift characters (as in the graphic type column definition used in a CREATE TABLE statement) nor the NULL indicator byte if the source column is nullable. If the length parameter is omitted, the default output field size is the length defined on the corresponding table column plus two bytes (shift-out and shift-in characters).

If the specified *length* is larger than the size of the data, the field is padded by the default DBCS padding character on the right.

TRUNCATE
The optional TRUNCATE keyword indicates that a graphic character string will be truncated from the right by the DBCS characters, if the data does not fit in the available space for the field in the output records. Without TRUNCATE, an error will occur when the output field size is too small for the data. An error can also occur with the TRUNCATE option if the available space is less than four bytes (four bytes will be the minimum size for a GRAPHIC EXTERNAL field; SO, one DBCS, and SI); or less than five bytes if the field is nullable (the four bytes plus the NULL indicator byte).

VARGRAPHIC
Specifies that the output field is of the varying length graphic type. A two byte binary length field is prepended to the actual data field. If the table column is nullable, a NULL indicator byte is placed before the length field.

(length)
Specifies the maximum length of the actual data field in the number of

DBCS characters. If you also specified NOPAD, it indicates the maximum allowable space for the data in the output records; otherwise, the space of the specified length is reserved for the data.

If the length parameter is omitted, the default is the smaller of 127 and the maximum length defined on the source table column.

STRIP
With this option, DBCS blanks (the default) or the specified characters are removed from the unloaded data and the VARGRAPHIC length field is adjusted to the length of the stripped data (the number of DBCS characters).

The effect of the STRIP option is the same as the SQL STRIP scalar function. For details, see Chapter 5 of *DB2 SQL Reference*.

BOTH
Indicates that occurrences of blank or the specified character are removed from both ends. The **default** is **BOTH**.

TRAILING
Indicates that occurrences of blank or the specified character are removed from the end.

LEADING
Indicates that occurrences of blank or the specified character are removed from the beginning.

X'strip-char'
Specifies a DBCS character to be stripped in the hexadecimal format, **X**'*hhhh*', where *hhhh* is four hexadecimal characters representing a DBCS character. If this operand is omitted, the default is a DBCS blank in the output encoding scheme (for example, X'4040' for the EBCDIC-encoded output, X'8140' for CCSID 301).

The strip operation is applied after the character code conversion, if the output character encoding scheme is different from the one defined on the source data. Therefore, if you specify a strip character, it must be in the encoding scheme used for the output.

TRUNCATE
The optional TRUNCATE keyword indicates that a graphic character string (encoded for output) will be truncated from the right, if the data does not fit in the available space for the field in the output records. Truncation occurs at a DBCS character boundary. Without TRUNCATE, an error will occur when the output field size is too small for the data.

SMALLINT
Specifies that the output field is a two byte binary integer (a negative number is in two's complement notation). If the external format is desired, use **INTEGER EXTERNAL**.

If the source data type is INTEGER, DECIMAL, or FLOAT (either four byte or eight byte format), an error will occur when the data is greater than 32 767 or less than -32 768.

A SMALLINT output field requires two bytes and the *length* option is not available.

INTEGER

Specifies that the output field is a four byte binary integer (a negative number is in two's complement notation).

If the original data type is DECIMAL, or FLOAT (either four byte or eight byte format), an error will occur when the original data is greater than 2 147 483 647 or less than -2 147 483 648.

An INTEGER output field requires four bytes and the *length* option is not available.

INTEGER EXTERNAL

Specifies that the output field contains a character string representing an integer number.

(length)

The optional *length* parameter indicates the size of the output data in bytes, including a space for the sign character. When the *length* is given and the character notation does not fit in the space, an error will occur. If the output field size is not specified, the default is 11 characters (including a space for the sign).

If the value is negative, a minus sign precedes the numeric digits. If the output field size is larger than the length of the data, the output data is left justified and blanks will be padded on the right.

If the source data type is DECIMAL, or FLOAT (either four byte or eight byte format), an error will occur when the original data is greater than 2 147 483 647 or less than -2 147 483 648.

DECIMAL

Specifies that the output data is a number represented by one of the decimal formats described below. If you specify the keyword DECIMAL by itself, packed decimal format is assumed.

PACKED

Specifies that the output data is a number represented by the packed decimal format. You can use **DEC** or **DEC PACKED** as an abbreviated form of the keyword.

The packed decimal representation of a number is of the form *ddd...ds*, where *d* is a decimal digit represented by four bits, and *s* is a four bit sign character (hexadecimal A, C, E, or F for a positive number, and hexadecimal B or D for a negative number).

length

Specifies the number of digits (not including the sign digit) placed in the output field. The length must be between one and 31. If the length is odd, the size of the output data field is (length+1)/2 bytes; if even, (length/2)+1 byte.

If the source data type is DECIMAL and the *length* parameter is omitted, the default length is determined by the column attribute defined on the table. Otherwise, the default length is 31 digits (16 bytes).

scale

Specifies the number of digits to the right of the decimal point (note that a decimal point is not included in the output field in this case). The number must be an integer greater than or equal to zero and must be less than or equal to the length.

The default depends on the column attribute defined on the table. If the source data type is DECIMAL, then the defined *scale* value is taken as the default value; otherwise, the **default** is **0**.

If you specify the output field size as less than the length of the data, an error will occur. If the specified field size is greater than the length of data, X'0' will be padded on the left.

ZONED

Specifies that the output data is a number represented by the zoned decimal format. You can use DEC ZONED as an abbreviated form of the keyword.

The zoned decimal representation of a number is of the form *znznzn...z/sn*, where *n* denotes a four bit decimal digit (called the numeric bits); *z* is the digit's zone (left four bits of a byte); *s* is the right most operand that can be a zone (*z*) or can be a sign value (hexadecimal A, C, E, or F for a positive number, and hexadecimal B or D for a negative number).

length

Specifies the number of bytes (i.e., the number of decimal digits) placed in the output field. The length must be between one and 31.

If the source data type is DECIMAL and the *length* parameter is omitted, the default length is determined by the column attribute defined on the table. Otherwise, the default length is 31 bytes.

scale

Specifies the number of digits to the right of the decimal point (note that a decimal point is not included in the output field in this case). The number must be an integer greater than or equal to zero and must be less than or equal to the length.

The default depends on the column attribute defined on the table. If the source data type is DECIMAL, then the defined *scale* value is taken as the default value; otherwise, the **default** is **0**.

If you specify the output field size as less than the length of the data, an error will occur. If the specified field size is greater than the length of data, X'F0' will be padded on the left.

EXTERNAL

Specifies that the output data is a character string representing a number in the form of ±*dd...d.ddd...d* where *d* is a numeric character 0-9 (the plus sign for a positive value is omitted).

length

Specifies the overall length of the output data (the number of characters including a sign, and a decimal point if *scale* is specified).

If the source data type is DECIMAL and the *length* parameter is omitted, the default length is determined by the column attribute defined on the table. Otherwise, the default length is 33 (31 numeric digits, a sign, and a decimal point). The minimum value of length is 3 to accommodate the sign, one digit, and the decimal point.

scale

Specifies the number of digits to the right of the decimal point. The

number must be an integer greater than or equal to zero and must be less than or equal to *length–2* (to allow the sign character and the decimal point).

The default depends on the column attribute defined on the table. If the source data type is DECIMAL and the *length* parameter is omitted, the default scale is determined by the column attribute defined on the table. Otherwise, the **default** is **0**.

An error will occur if the character representation of a value does not fit in the given or default field size (precision). If the source data type is floating point and a data item is too small for the precision defined by *scale*, the value of zero will be returned (not an error).

FLOAT*(length)*

Specifies that the output data is a binary floating point (32 bit or single precision FLOAT if the *length* is between one and 21 inclusive; 64 bit or double precision FLOAT if the *length* is between 22 and 53 inclusive). If the *length* parameter is omitted, the 64 bit format is assumed (output field size will be eight bytes). Note that the *length* parameter for the FLOAT type does not represent the field size in bytes.

The format of the binary floating point output is controlled by the global FLOAT option. The default is S/390 format (Hexadecimal Floating Point or HFP). If you specify FLOAT(IEEE), all the binary floating point output will be in IEEE format (Binary Floating Point or BFP). When you specify FLOAT(IEEE) and the source data type DOUBLE is unloaded as REAL, an error will occur if the source data can not be expressed by the IEEE (BFP) 32-bit notation.

EXTERNAL *(length)*

Specifies that the output data is a number represented by a character string in floating point notation, ±*d.ddd...dddE*±*nn*, where *d* is a numeric character (0-9) for the significant digits; *nn* after the character *E* and the sign consists of two numeric characters for the exponent.

length

Specifies the total field length in bytes including the first sign character, the decimal point, the *E* character, the second sign character and the two digit exponent. If the number of characters in the result is less than the specified or the default length, the result is padded to the right with blanks. The length, if specified, must be greater than or equal to eight.

The default output field size is 14 if the source data type is the 32 bit FLOAT, 24 otherwise.

A FLOAT EXTERNAL output field requires a space of at least seven characters in the output record to accommodate the minimal floating point notation. Otherwise, an error will occur.

REAL The keyword **REAL** can be used for 32 bit floating point output. If REAL is used, the *length* parameter must not be specified.

DOUBLE

The keyword **DOUBLE** can be used for 64 bit floating point output. If DOUBLE is used, the *length* parameter must not be specified.

DATE EXTERNAL

Specifies that the output field is for a character string representation of a date. The output format of date depends on the DB2 installation.

(length)
Specifies the size of the data field in bytes in the output record. A DATE EXTERNAL field requires a space of at least 10 characters. If the space is not available, an error will occur. If the specified *length* is larger than the size of the data, blanks will be padded on the right.

TIME EXTERNAL
Specifies that the output field is for a character string representation of a time. The output format of time depends on the DB2 installation.

(length)
Specifies the size of the data field in bytes in the output record. A TIME EXTERNAL field requires a space of at least eight characters. If the space is not available, a conversion error will occur. If the specified *length* is larger than the size of the data, blanks will be padded on the right.

TIMESTAMP EXTERNAL
Specifies that the output field is for a character string representation of a timestamp. The output format of timestamp will be

- *yyyy-mm-dd-hh.mm.ss* if the *length* is not specified (default) or the specified *length* is greater than or equal to 19 and less than 26.
- *yyyy-mm-dd-hh.mm.ss.nnnnnn* if the *length* is greater than or equal to 26.

(length)
Specifies the size of the data field in bytes in the output record. A TIMESTAMP EXTERNAL field requires a space of at least 19 characters. If the space is not available, an error will occur. The *length* parameter, if specified, will determine the output format of the TIMESTAMP (see above). If the specified *length* is larger than the size of the data, the field is padded by the default padding character on the right.

CONSTANT
Specifies that the output records will have an extra field containing a constant value. The field name associated with the CONSTANT keyword must not coincide with a table column name (the field name is for clarification purposes only). A CONSTANT field always has a fixed length equal to the length of the given string.

'string'
Specifies the character string inserted in the output records at the specified or at the default position. A string is the required operand of the CONSTANT option. If the given string is in the form *'string'*, it is assumed to be an EBCDIC SBCS string. The output string for a CONSTANT field is, however, in the specified or the default encoding scheme, that is, if the encoding scheme used for output is not EBCDIC, the SBCS CCSID conversion is applied to the given string before it is placed in output records.

X'hex-string'

The hexadecimal form, X'*hex-string*' can also be used to represent a string constant. If it is desired to specify a CONSTANT string value in an encoding scheme other than SBCS EBCDIC, use the hexadecimal form. No CCSID conversion is performed if the hexadecimal form is used.

For a CONSTANT field, no other field selection list options should be specified.

If a CONSTANT field is inserted, it will not be included in the generated LOAD statement (the LOAD statement will be generated so that the CONSTANT field is skipped).

ROWID

The field type ROWID can be specified if and only if the column to be unloaded is of type ROWID. The keyword is provided for consistency purposes.

ROWID fields have varying length and a two byte binary length field will be prepended to the actual data field.

For the ROWID type, no data conversion nor truncation will be applied. If the output field size is too small to unload ROWID data, an error will occur.

If the source is an image copy and a ROWID column is selected, and if the pageset header page is missing in the specified data set, the UNLOAD utility terminates with the error message DSNU1228I. This can happen when an image copy data set of DSNUM greater than one for a non-partitioned table space defined on multiple data sets is specified as the source of unloading.

BLOB Indicates that the column is unloaded as a binary large object (BLOB). No data conversion will be applied to the field.

When you specify the BLOB field type, a four byte binary length field is placed in the output record prior to the actual data field. If the source table column is nullable, a NULL indicator byte is placed before the length field.

(length)

Specifies the maximum length of the actual data field in bytes. If you specify NOPAD, it indicates the maximum allowable space for the data in the output records; otherwise, the space of the specified length is reserved for the data.

If the length parameter is omitted, the default is the maximum length defined on the source table column.

TRUNCATE

The optional TRUNCATE keyword indicates that a BLOB string will be truncated from the right, if the data does not fit in the available space for the field in the output record. For BLOB data, truncation occurs at a byte boundary. Without TRUNCATE, an error will occur when the output field size is too small for the data.

CLOB Indicates that the column is unloaded as a character large object (CLOB).

When you specify the CLOB field type, a four byte binary length field is placed in the output record prior to the actual data field. If the source table column is nullable, a NULL indicator byte is placed before the length field.

If you specified the EBCDIC, ASCII, UNICODE, or CCSID options, the output data is encoded in the CCSID corresponding to the specified option, depending on the subtype of the source data (SBCS or MIXED). No conversion will be applied if the subtype is BIT.

(length)

Specifies the maximum length of the actual data field in bytes. If you

specify NOPAD, it indicates the maximum allowable space for the data in the output records; otherwise, the space of the specified length is reserved for the data.

If the length parameter is omitted, the default is the maximum length defined on the source table column.

TRUNCATE
The optional TRUNCATE keyword indicates that a CLOB string (encoded for output) will be truncated from the right, if the data does not fit in the available space for the field in the output record. For a CLOB data, truncation occurs at a character boundary. See "Specifying TRUNCATE and STRIP options for output data" on page 516 for the truncation rules used in the UNLOAD utility. Without TRUNCATE, an error will occur when the output field size is too small for the data.

DBCLOB (double-byte character large object)
Indicates that the column is unloaded as a double-byte character large object (DBCLOB).

If you specify the DBCLOB field type, a four byte binary length field is placed in the output record prior to the actual data field. If the source table column is nullable, a NULL indicator byte is placed before the length field.

If you specified the EBCDIC, ASCII, UNICODE, or CCSID options, the output data is encoded in the CCSID corresponding to the specified option (DBCS CCSID is used).

(length)
Specifies the maximum length of the actual data field in the number of DBCS characters. If you specify NOPAD, it indicates the maximum allowable space for the data in the output records; otherwise, the space of the specified length is reserved for the data.

If the length parameter is omitted, the default is the maximum length defined on the source table column.

TRUNCATE
The optional TRUNCATE keyword indicates that a DBCS string (encoded for output) will be truncated from the right, if the data does not fit in the available space for the field in the output record. For a DBCLOB data, truncation occurs at a character (DBCS) boundary. See "Specifying TRUNCATE and STRIP options for output data" on page 516 for the truncation rules used in the UNLOAD utility. Without TRUNCATE, an error will occur when the output field size is too small for the data.

WHEN
The WHEN clause tells which records in the table space are to be unloaded. If there is no WHEN clause specified for a table in the table space, all of the records are unloaded.

The option following WHEN describes the conditions for unloading records from a table.

selection condition
A *selection condition* specifies a condition that is true, false, or unknown about a given row. When the condition is true, the row qualifies for UNLOAD. When the condition is false or unknown, the row does not qualify.

The result of a selection condition is derived by application of the specified *logical operators* (AND, OR) to the result of each specified predicate. If logical operators are not specified, the result of the selection condition is the result of the specified predicate.

Selection conditions within parentheses are evaluated first. If the order of evaluation is not specified by parentheses, AND is applied before OR.

predicate
A *predicate* specifies a condition that is true, false, or unknown about a row.

labeled duration expression
A *labeled duration expression* specifies an expression that begins with special register CURRENT DATE or special register CURRENT TIMESTAMP (the forms CURRENT_DATE and CURRENT_TIMESTAMP are also acceptable) and optionally arithmetic operations of addition or subtraction expressed by a number followed by one of the seven duration keywords: YEARS, MONTHS, DAYS, HOURS, MINUTES, SECONDS, or MICROSECONDS. (The singular form of these keywords is also acceptable: YEAR, MONTH, DAY, HOUR, MINUTE, SECOND, and MICROSECOND.)

Utilities always evaluate a *labeled duration expression* as a timestamp and implicitly convert to a date if the comparison is with a date column.

Incrementing and decrementing CURRENT DATE: The result of adding a duration to a date, or of subtracting a duration from a date, is itself a date. (For the purposes of this operation, a month denotes the equivalent of a calendar page. Adding months to a date, then, is like turning the pages of a calendar, starting with the page on which the date appears.) The result must fall between the dates January 1, 0001 and December 31, 9999 inclusive. If a duration of years is added or subtracted, only the year portion of the date is affected. The month is unchanged, as is the day unless the result would be February 29 of a non-leap-year. In this situation, the day portion of the result is set to 28.

Similarly, if a duration of months is added or subtracted, only months and, if necessary, years are affected. The day portion of the date is unchanged unless the result would be invalid (September 31, for example). In this case the day is set to the last day of the month.

Adding or subtracting a duration of days will affect the day portion of the date, and potentially the month and year.

Date durations, whether positive or negative, can also be added to and subtracted from dates. As with labeled durations, the result is a valid date.

When a positive date duration is added to a date, or a negative date duration is subtracted from a date, the date is incremented by the specified number of years, months, and days.

When a positive date duration is subtracted from a date, or a negative date duration is added to a date, the date is decremented by the specified number of days, months, and years.

Adding a month to a date gives the same day one month later unless that day does not exist in the later month. In that case, the day in the result is set to the last day of the later month. For example, January 28

plus one month gives February 28; one month added to January 29, 30, or 31 results in either February 28 or, for a leap year, February 29. If one or more months is added to a given date and then the same number of months is subtracted from the result, the final date is not necessarily the same as the original date.

The order in which labeled date durations are added to and subtracted from dates can affect the results. When you add labeled date durations to a date, specify them in the order of YEARS + MONTHS + DAYS. When you subtract labeled date durations from a date, specify them in the order of DAYS - MONTHS - YEARS. For example, to add one year and one day to a date, specify:

```
CURRENT DATE + 1 YEAR + 1 DAY
```

To subtract one year, one month, and one day from a date, specify:

```
CURRENT DATE - 1 DAY - 1 MONTH - 1 YEAR
```

Incrementing and decrementing timestamps: The result of adding a duration to a timestamp, or of subtracting a duration from a timestamp, is itself a timestamp. Date and time arithmetic is performed as previously defined, except that an overflow or underflow of hours is carried into the date part of the result, which must be within the range of valid dates.

basic predicate

A *basic predicate* compares a column with a constant. If the value of the column is null, the result of the predicate is unknown. Otherwise, the result of the predicate is true or false.

column = constant	column is equal to constant or labeled duration expression.
column < > constant	column is not equal to constant or labeled duration expression.
column > constant	column is greater than constant or labeled duration expression.
column < constant	column is less than constant or labeled duration expression.
column > = constant	column is greater than or equal to constant or labeled duration expression.
column < = constant	column is less than or equal to constant or labeled duration expression.

Note: The following alternative comparison operators are available:
!= or **¬=** for not equal.
!> or **¬>** for not greater than.
!< or **¬<** for not less than.

The symbol ¬ representing “not” is supported for compatibility purposes. Use **!** where possible.

BETWEEN predicate

The BETWEEN predicate determines whether a given value lies between two other given values specified in ascending order. Each of the predicate’s two forms has an equivalent search condition, as shown below:

The predicate:	*column* BETWEEN *value1* AND *value2*
is equivalent to:	(*column* >= *value1* AND *column* <= *value2*)
The predicate:	*column* NOT BETWEEN *value1* AND *value2*
is equivalent to:	NOT(*column* BETWEEN *value1* AND *value2*)
and therefore also to:	(*column* < *value1* OR *column* > *value2*)

The values can be constants or labeled duration expressions.

For example, the following predicate is true for any row when salary is greater than or equal 10000 and less than or equal to 20000 :

```
SALARY BETWEEN 10000 AND 20000
```

IN predicate

The IN predicate compares a value with a set of values. In the IN predicate, the second operand is a set of one or more values specified by constants.

The predicate:	*value1* IN (*value1*, *value2*,..., *valuen*)
is equivalent to:	(*value1* = *value2* OR ... OR *value1* = *valuen*)
The predicate:	*value1* NOT IN (*value1*, *value2*,..., *valuen*)
is equivalent to:	(*value1* ¬= *value2* AND ... AND *value1* ¬= *valuen*)

For example, the following predicate is true for any row whose employee is in department D11, B01, or C01:

```
WORKDEPT IN ('D11', 'B01', 'C01')
```

LIKE predicate

The *LIKE predicate* qualifies strings that have a certain pattern. The pattern is specified by a string in which the underscore and percent sign characters have special meanings.

Let *x* denote the column to be tested and *y* the pattern in the string constant.

The following rules apply to predicates of the form ″x LIKE y...″. If NOT is specified, the result is reversed.

- When *x* and *y* are both neither empty nor null, the result of the predicate is true if *x* matches the pattern in *y* and false if *x* does not match the pattern in *y*. Matching the pattern is described below.
- When *x* or *y* is null, the result of the predicate is unknown.
- When *y* is empty and *x* is not, the result of the predicate is false.
- When *x* is empty and *y* is not, the result of the predicate is false unless *y* consists only of one or more percent signs.
- When *x* and *y* are both empty, the result of the predicate is true.

The pattern string and the string to be tested must be of the same type, that is, both *x* and *y* must be character strings or both *x* and *y* must be graphic strings. When *x* and *y* are graphic strings, a character is a DBCS character. When *x* and *y* are character strings and *x* is not mixed data, a character is an SBCS character and *y* is interpreted as SBCS data regardless of its subtype. The rules for mixed data patterns are described on page 507.

Within the pattern, a percent sign or underscore can have a special meaning, or it can represent the literal occurrence of a percent sign or underscore. To have its literal meaning, it must be preceded by an *escape character*. If it is not preceded by an escape character, it has its special meaning.

The ESCAPE clause designates a single character. That character, and only that character, can be used multiple times within the pattern as an escape character. When the ESCAPE clause is omitted, no character serves as an escape character, so that percent signs and underscores in the pattern always have their special meanings.

The following rules apply to the use of the ESCAPE clause:

- The ESCAPE clause cannot be used if *x* is mixed data.
- If *x* is a character string, the data type of the string constant must be character string. If *x* is a graphic string, the data type of the string constant must be graphic string. In both cases, the length of the string constant must be 1.
- The pattern must not contain the escape character except when followed by the escape character, '%' or '_'. For example, if '+' is the escape character, any occurrences of '+' other than '++', '+_', or '+%' in the pattern is an error.

When the pattern does not include escape characters, a simple description of its meaning is:

- The underscore sign (_) represents a single arbitrary character.
- The percent sign (%) represents a string of zero or more arbitrary characters.
- Any other character represents a single occurrence of itself.

More rigorous description of strings and patterns

The string *y* is interpreted as a sequence of the minimum number of substring specifiers such that each character of *y* is part of exactly one substring specifier. A substring specifier is an underscore, a percent sign, or any non-empty sequence of characters other than an underscore or percent sign.

The string *x* matches the pattern *y* if a partitioning of *x* into substrings exists, such that:

- A substring of *x* is a sequence of zero or more contiguous characters and each character of *x* is part of exactly one substring.
- If the *n*th substring specifier is an underscore, the *n*th substring of *x* is any single character.
- If the *n*th substring specifier is a percent sign, the *n*th substring of *x* is any sequence of zero or more characters.
- If the *n*th substring specifier is neither an underscore nor a percent sign, the *n*th substring of *x* is equal to that substring specifier and has the same length as that substring specifier.
- The number of substrings of *x* is the same as the number of substring specifiers.

When escape characters are present in the pattern string, an underscore, percent sign, or escape character represents a single occurrence of itself if and only if it is preceded by an odd number of successive escape characters.

Mixed data patterns: If *x* is mixed data, the pattern is assumed to be mixed data, and its special characters are interpreted as follows:

- A single-byte underscore refers to one single-byte character; a double-byte underscore refers to one double-byte character.
- A percent sign, either single-byte or double-byte, refers to any number of characters of any type, either single-byte or double-byte.
- Redundant shift bytes in *x* or *y* are ignored.

NULL predicate

The *NULL predicate* tests for null values.

If the value of the column is null, the result is true. If the value is not null, the result is false. If NOT is specified, the result is reversed.

Instructions for running UNLOAD

To run UNLOAD, you must:

1. Prepare the necessary data sets, as described in “Data sets used by UNLOAD” on page 508.
2. Create JCL statements, by using one of the methods described in “Chapter 3. Invoking DB2 online utilities” on page 23. (For examples of JCL for UNLOAD, see “Sample control statements” on page 520.)
3. Prepare a utility control statement, specifying the options for the tasks you want to perform, as described in “Instructions for specific tasks” on page 508. (For a

complete description of the syntax and options for UNLOAD, see "Syntax and options of the control statement" on page 477.)

4. Check the compatibility table in "Concurrency and compatibility" on page 518 if you want to run other jobs concurrently on the same target objects.
5. Plan for restart if the UNLOAD job doesn't complete, as described in "Terminating or restarting UNLOAD" on page 518.
6. Run UNLOAD.

See "Chapter 3. Invoking DB2 online utilities" on page 23 for an explanation of ways to execute DB2 utilities.

Data sets used by UNLOAD

Table 99 describes the data sets used by UNLOAD. Include statements in your JCL for each required data set, and any optional data sets you want to use.

Table 99. Data sets used by UNLOAD

Data Set	Description	Required?
SYSIN	Input data set containing the utility control statement.	Yes
SYSPRINT	Output data set for messages.	Yes
Input data set	The input data set containing the data to be loaded. Its name is identified by the DD statement specified by the INDDN option. The default name is SYSREC. It must be a sequential data set that is readable by BSAM.	Yes
SYSPUNCH	One or more work data sets to contain the generated LOAD statements for subsequently reloading the data. The default DD name is PUNCHDDN.	No[1]
Unload data set	One or more work data set to contain the unloaded table rows. The default DD name is UNLDDN.	Yes

Note:
[1] Required if you request LOAD statement generation by specifying PUNCHDDN in the utility control statement.

The following object is named in the utility control statement and does not require a DD card in the JCL:

Table space
The name of the table space to be unloaded. It is named in the UNLOAD control statement and is accessed through the DB2 catalog. If you want to unload only one partition of a table space, you must specify the PART option in the control statement.

Instructions for specific tasks

The following tasks are described here:

"Unloading partitions" on page 509
"Selecting tables and rows to unload" on page 509
"Selecting and ordering columns to unload" on page 509
"Using inline image copies with UNLOAD" on page 509
"Converting data with the UNLOAD utility" on page 511

"Specifying output field types" on page 511
"Specifying output field positioning and size" on page 513
"Determining the layout of output fields" on page 514
"Specifying TRUNCATE and STRIP options for output data" on page 516
"Generating LOAD statements" on page 517
"Unloading compressed data" on page 517
"Interpreting field specification errors" on page 518

Unloading partitions

If the source table space is partitioned, use one of the following mutually exclusive methods to select the partitions to unload:

- Use the LIST keyword with a LISTDEF that contains PARTLEVEL specifications. Partitions can be either included or excluded by the use of the INCLUDE and the EXCLUDE features of LISTDEF.
- Specify the PART keyword to select a single partition or a range of partitions.

With either method, the unloaded data can be stored in a single data set for all selected partitions or in one data set for each selected partition. If you want to unload to a single output data set, specify a DD name to UNLDDN. If you want to unload into multiple output data sets, specify a template name that is associated with the partitions. You can process multiple partitions in parallel if the TEMPLATE definition contains the partition as a variable, for example `&PA`.

You can not specify multiple output data sets with the FROMCOPY or the FROMCOPYDDN option.

Selecting tables and rows to unload

If a table space contains multiple tables, you can selected specific tables to unload using the FROM TABLE specification clauses. If you specify one or more FROM TABLE clauses for a table space, only the qualified rows from the specified tables are unloaded. You can specify a maximum of one FROM TABLE clause per table. If you do not specify at least one FROM TABLE clause, the rows from all the tables in the table space will be unloaded.

Within a FROM TABLE clause, you can specify one or more of the following criteria:

- Row and column selection criteria using the field specification list
- Row selection conditions using WHEN specification clause
- Row sampling specifications

Selecting and ordering columns to unload

Use a field specification list in a FROM TABLE clause to unload specified columns in the listed order. If you omit a field specification list, all the columns in the row will be unloaded in the order of the columns defined on the table.

You can specify a format conversion option for each field in the field specification list.

If you select a LOB column in a list of field specifications or select a LOB column by default (by omitting a list of field specifications), LOB data is materialized in the output. However, you can not select LOB columns from image copy data sets.

Using inline image copies with UNLOAD

In addition to unloading data from table spaces and partitions, you can also unload data from one or more image copy data sets. Unload rows from a single image copy data set by specifying the FROMCOPY option in the UNLOAD control statement. Specify the FROMCOPYDDN option to unload data from one or more

image copy data sets associated with the specified DDNAME. Use an image copy that contains the page set header page when you are unloading a ROWID column, otherwise the unload fails.

The source image copy data set must have been created by one of the following utilities:

- COPY
- LOAD inline image copy
- MERGECOPY
- REORG TABLESPACE inline image copy
- DSN1COPY

UNLOAD accepts full image copies, incremental image copies, and a copy of pieces as valid input sources.

The UNLOAD utility supports image copy data sets for a single table space. The table space name must be specified in the TABLESPACE option. The specified table space must exist when you run the UNLOAD utility (that is, the table space has not been dropped since the image copy was taken).

Use the FROMCOPYDDN option to concatenate copy of pieces under a DDNAME to form a single input data set image. When you use the FROMCOPYDDN option, concatenate the data sets in the order of the data set number; the first data set must be concatenated first. If the data sets are concatenated in the wrong order or different generations of image copies are concatenated (for example, the most recent image copy data sets and older image copies are intermixed), the results might be unpredictable.

It is possible to use the FROMCOPYDDN option to concatenate a full image copy and incremental image copies for a table space, a partition, or a piece, but duplicate rows will also be unloaded in this situation. Instead, consider using MERGECOPY to generate an updated full image copy as the input to the UNLOAD utility.

You can select specific rows and columns to unload just as you would for a table space. However, you can only unload rows containing LOB columns when the LOB columns are not included in a field specification list. If you use an image copy that does not contain the page set header page when unloading a ROWID column, the unload fails.

If you use the FROMCOPY or the FROMCOPYDDN option, you can only specify one output data set.

If an image copy was created by an inline copy operation (LOAD or REORG TABLESPACE), the image copy can contain duplicate pages. If duplicate pages exist, the UNLOAD utility issues a warning message, and all the qualified rows in duplicate pages will be unloaded into the output data set.

If a dropped table is specified to the FROM TABLE option, the UNLOAD utility terminates with return code 8. If you do not specify a FROM TABLE and if an image copy contains rows from dropped tables, UNLOAD ignores these rows. When you specify a copy of pieces (either full or incremental) of a segmented table space consisting of multiple data sets in the FROMCOPY option, and a mass delete was applied to a table in the table space before the copy was created, deleted rows will be unloaded if the space map pages indicating the mass delete are not included in

the data set corresponding to the specified copy. Where possible, use the FROMCOPYDDN option to concatenate the copy of pieces.

If an image copy contains a table to which ALTER ADD COLUMN was applied after the image copy was taken , the UNLOAD utility sets the system or user specified default value for the added column when the data is unloaded from such an image copy.

Converting data with the UNLOAD utility

You can convert one data type to another compatible data type using the UNLOAD utility.[5] For example, you can convert columns of a numeric type (SMALLINT, INTEGER, FLOAT, DOUBLE, REAL, and DECIMAL) from the DB2 internal format to the S/390 or an external format.

When you unload a floating point type column, you can specify the binary form of the output to either the S/390 format (hexadecimal floating point, or HFP), or the IEEE format (binary floating point, or BFP).

You can also convert a varying length column to a fixed length output field, with or without padding characters. In either case, unless you explicitly specify a fixed length data type for the field, the data itself is treated as a varying length data and a length field is prepended to the data.

For certain data types, you can unload data into fields with a smaller length using the TRUNCATE or STRIP options. In this situation, if a character code conversion is applied, the length of the data in bytes might change due to the code conversion. The truncation operation is applied after the code conversion.

You can perform character code conversion on a character type field, including converting numeric columns to the external format and the CLOB type. Be aware that when you apply a character code conversion for mixed data fields, the length of the result string in bytes can be shorter or longer than the length of the source string. Character type data always gets converted if you specified any of the character code conversion options (EBCDIC, ASCII, UNICODE or CCSID).

DATE, TIME, or TIMESTAMP column types will always be converted into the external formats based on the DATE, TIME, and TIMESTAMP formats of your installation.

For specific data type compatibility options, see “Specifying output field types”.

Specifying output field types

An output field can have a different data type from the one defined on a source table column as long as the data types are compatible. The UNLOAD utility follows the general rules and conventions of DB2 for OS/390 and z/OS on the data type attributes and the compatibility among the data types as described in Chapter 2 of *DB2 SQL Reference*.

If you specify a data type in the UNLOAD control statement, the field type information will be included in the generated LOAD utility statement. For specific data type compatibility information, refer to Table 100 on page 512, Table 101 on page 512 and Table 102 on page 512.

5. The source type is used for user-defined distinct types.

Table 100. Numeric data conversion

Input Data Types	Output Data Types				
	SMALLINT	INTEGER (External)	DECIMAL (External)	FLOAT (External)	DOUBLE or REAL
SMALLINT	Y	Y[2]	Y[2]	Y[2]	Y
INTEGER	Y[6]	Y[2]	Y[2]	Y[2]	Y
DECIMAL	Y[6]	Y[2, 6]	Y[2]	Y[2]	Y
FLOAT, DOUBLE, or REAL	Y[6]	Y[2, 6]	Y[2, 6]	Y[2]	Y

Table 101. Character data conversion

Input Data Types	Output Data Types							
	BLOB	CHAR	VARCHAR	CLOB	GRAPHIC	GRAPHIC EXTERNAL	VAR-GRAPHIC	DBCLOB
BLOB	Y	N	N	N	N	N	N	N
CLOB	N	Y[1, 3]	Y[1, 3]	Y	N	N	N	N
DBCLOB	N	N	N	N	Y[1, 3]	Y[1, 3, 7]	Y[1, 3]	Y[1]
CHAR	N	Y[1]	Y[1]	Y[1, 8]	N	N	N	N
VARCHAR or LONG VARCHAR	N	Y[1, 3]	Y[1]	Y[1, 8]	N	N	N	N
GRAPHIC	N	N	N	N	Y[1]	Y[1, 7]	Y[1]	Y[1]
VARGRAPHIC or LONG VARGRAPHIC	N	N	N	N	Y[1, 3]	Y[1, 3, 7]	Y[1]	Y[1]

Table 102. Time data conversion

Input Data Types	Output Data Types		
	DATE EXTERNAL	TIME EXTERNAL	TIMESTAMP EXTERNAL
DATE	Y[2]	N	Y[2, 4]
TIME	N	Y[2]	N
TIMESTAMP	Y[2, 5]	Y[2, 5]	Y[2]

Table 102. Time data conversion (continued)

Input Data Types	Output Data Types		
	DATE EXTERNAL	TIME EXTERNAL	TIMESTAMP EXTERNAL

Notes:

1. Subject to the CCSID conversion, if specified.
2. Subject to the CCSID conversion, if specified (EXTERNAL case). See the description of the CCSID option.
3. Results in an error if the field length is too small for the data unless you specified the TRUNCATE option. Note that a LOB has a four byte length field; any other varying length type has a two byte length field.
4. Zeros in the time portion.
5. DATE or TIME portion of the timestamp.
6. Potential overflow (conversion error).
7. Only in the EBCDIC output mode.
8. Not applicable to BIT subtype data.

Specifying output field positioning and size

Output data is always placed by default in an output record in the order of the defined columns over the selected tables. You can choose to specify the order of the output fields using a list of field specifications.

Use the POSITION option to specify field position in the output records. You can also specify the size of the output data field using the *length* parameter for a particular data type. Note that the *length* parameter must indicate the size of the actual data field, and while the *start* parameter of the POSITION option indicates the starting position of a field including the NULL indicator byte (if the field is nullable) and the length field (if the field is varying length).

Using the POSITION parameter, the length parameter, or both can restrict the size of the data field in the output records. Use care when specifying the POSITION and length parameters, especially for nullable fields and varying length fields. If there is a conflict between the length parameter and the size of the field in the output record specified by the POSITION parameters, DB2 issues an error message and the UNLOAD utility terminates. If an error occurs, the count of the number of records in error is incremented. See the description of the MAXERR option on page 484 for more information.

If you specify a *length* parameter for a varying length field and you also specify the NOPAD option, *length* indicates the maximum length of data to be unloaded. Without the NOPAD option, UNLOAD reserves a space of the given *length* instead of the maximum data size.

If you explicitly specify *start* parameters for certain fields, they must be listed in ascending order in the field selection list. Note that unless you specified HEADER NONE for the table, a fixed length record header is placed at the beginning of each record for the table, and the *start* parameter must not overlap the record header area.

The TRUNCATE option is available for certain output field types. See "FROM-TABLE-spec" on page 485 and "Specifying TRUNCATE and STRIP options for output data" on page 516 for more information. For the output field types where the TRUNCATE option is not applicable, enough space must be provided in the

output record for each field. The output field layouts are summarized in "Determining the layout of output fields".

For errors that can occur at the record level due to the field specifications, see "Interpreting field specification errors" on page 518.

Determining the layout of output fields

Figure 27 shows a layout diagram of a fixed-length field (NOT NULL).

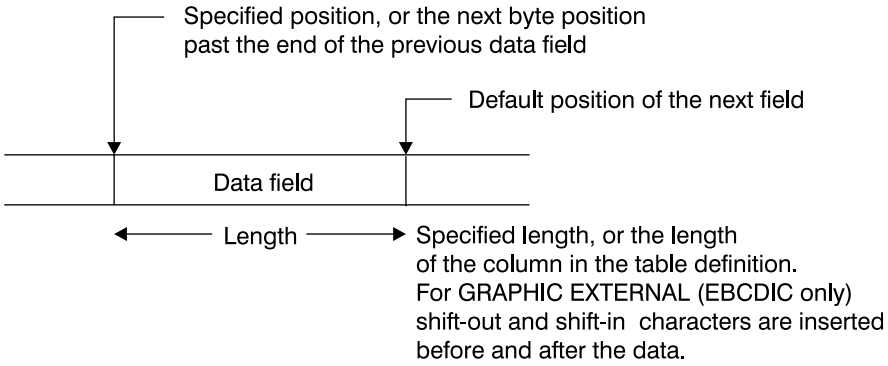

Figure 27. Layout diagram of a fixed-length field (NOT NULL)

Figure 28 shows a layout diagram of a nullable fixed-length field.

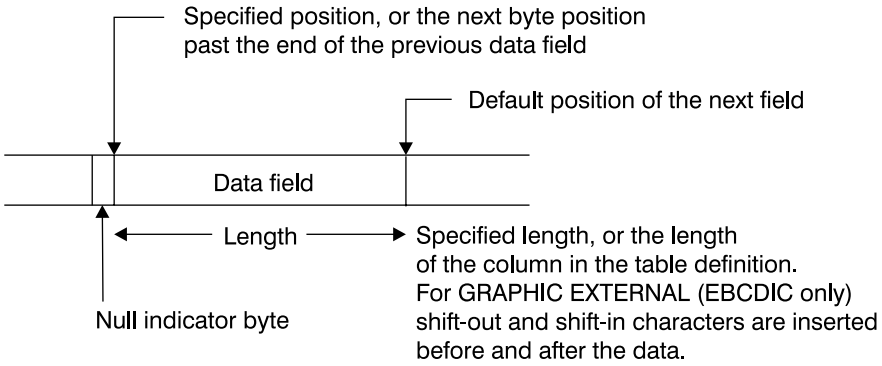

Figure 28. Layout diagram of a nullable fixed-length field

Figure 29 on page 515 shows a layout diagram of a varying-length field (NOT NULL) with the NOPAD option specified.

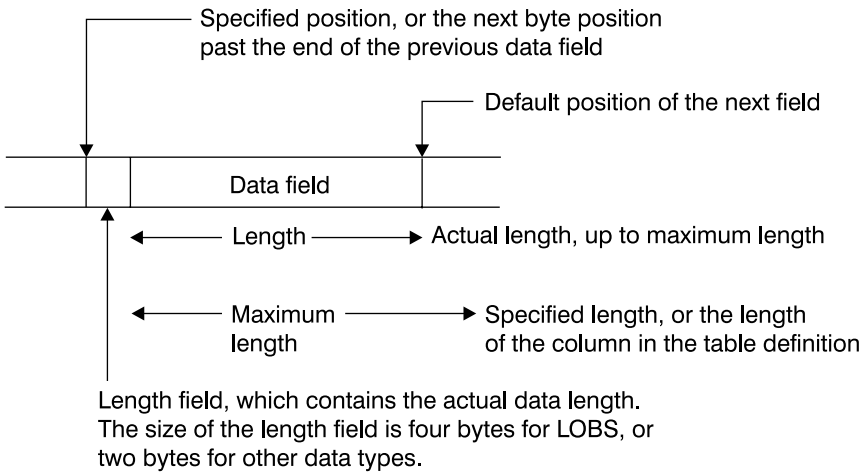

Figure 29. Layout diagram of a varying-length field (NOT NULL) with NOPAD

Figure 30 shows a layout diagram of a varying-length field (NOT NULL) without the NOPAD option specified.

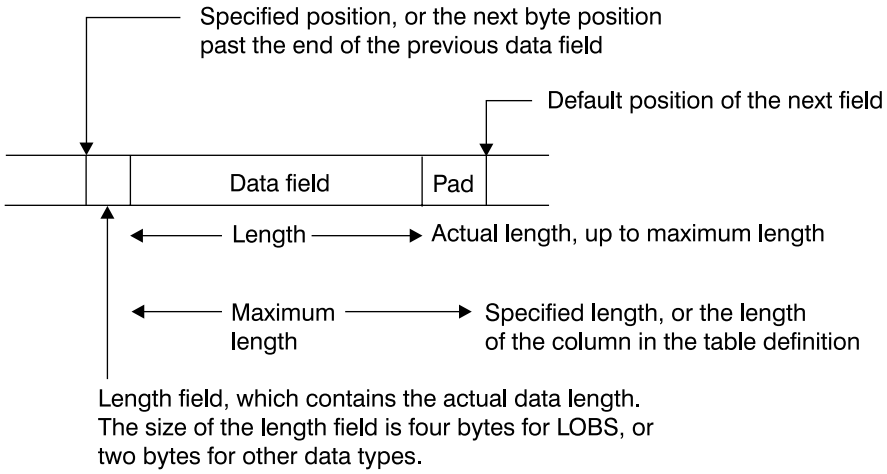

Figure 30. Layout diagram of a varying-length field (NOT NULL) without NOPAD

Figure 31 on page 516 shows a layout diagram of a nullable varying-length field with the NOPAD option specified.

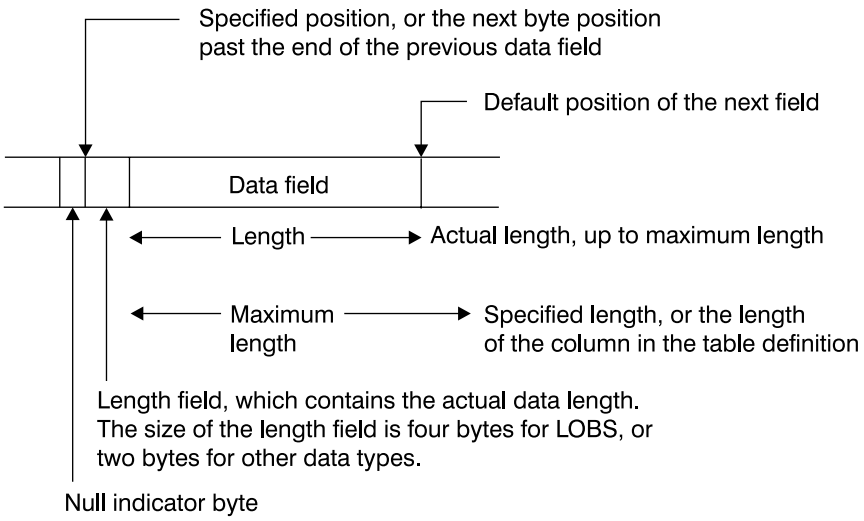

Figure 31. Layout diagram of a nullable varying-length field with NOPAD

Figure 32 shows a layout diagram of a nullable varying-length field without the NOPAD option specified.

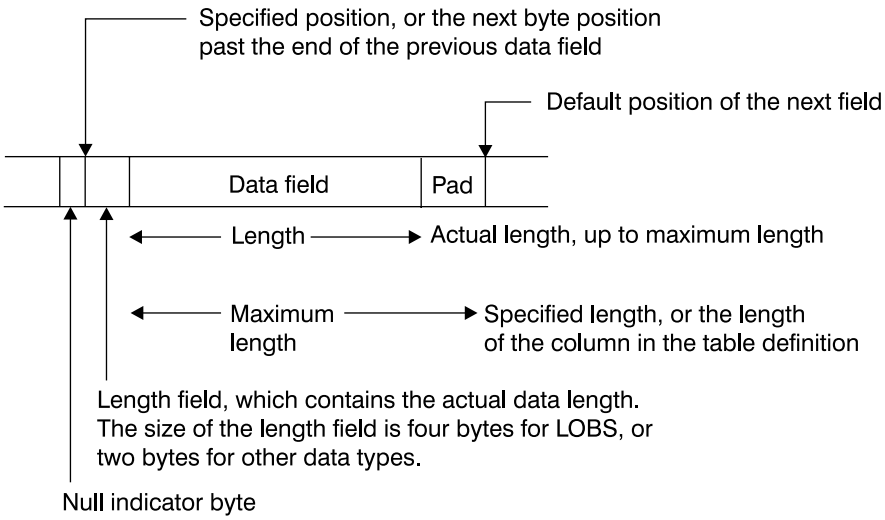

Figure 32. Layout diagram of a nullable varying-length field without NOPAD

Specifying TRUNCATE and STRIP options for output data

For certain types of data, it is possible to unload them into output fields shorter than the length of the output data. This data truncation occurs only when you explicitly specify the TRUNCATE option. Any CCSID conversion is applied first, then truncation is applied to encoded data for output.

For bit strings, truncation occurs at a byte boundary. For character type data, truncation occurs at a character boundary (a multi-byte character will not be split). If a MIXED character type data is truncated in an output field of fixed size, the truncated string can be shorter than the specified field size. In this case, blanks in the output CCSID will be padded to the right. If the output data is in EBCDIC for a MIXED character type field, truncation preserves the SO (shift out) and the SI (shift in) characters around a DBCS substring.

The TRUNCATE option of the UNLOAD utility is intended to truncate string data, and has a different purpose than the SQL TRUNCATE scalar function.

For VARCHAR and VARGRAPHIC output fields, in addition to the TRUNCATE option, the STRIP option is provided to remove the specified characters, or the leading blanks, the trailing blanks, or both. The STRIP operation is applied on the encoded data for output. If both the TRUNCATE and STRIP options are specified, the truncation operation is applied first, then STRIP is applied. For example, if both the TRUNCATE and STRIP options are specified for a VARCHAR(5) output field, see the example of output in Table 103. In the following table, an underscore represents a character to be stripped. Note that in all cases, the source string is first truncated to '_ABC_' (a five character string to fit in the VARCHAR(5) field), then the strip operation is applied.

Table 103. Results of specifying both TRUNCATE and STRIP options for UNLOAD

Source string	Output string	*length* specified	STRIP option specified
'_ABC_DEF'	'ABC'	3	**STRIP BOTH**
'_ABC_DEF'	'ABC_'	4	**STRIP LEADING**
'_ABC_DEF'	'_ABC'	4	**STRIP TRAILING**

Generating LOAD statements

To enable reloading the unloaded data into either the original table or different tables, a LOAD utility statement will be generated and written to the SYSPUNCH DDNAME or to the DDNAME specified by PUNCHDDN.

The generated LOAD statement includes WHEN and INTO TABLE specifications that identify the table where the rows are to be reloaded, unless the HEADER NONE option was specified in the UNLOAD control statement. You will need to edit the generated LOAD statement if you intend to load the UNLOAD output data into different tables than the original ones.

If multiple table spaces are to be unloaded and the generation of LOAD statements are desired, you must specify a physically distinct data set for each table space to PUNCHDDN using a TEMPLATE containing the table space as a variable (&TS.).

If PUNCHDDN is not specified and the SYSPUNCH DDNAME does not exist, the LOAD statement will not be generated.

Unloading compressed data

Compressed rows from an image copy data set can be unloaded only when the dictionary for decompression has been retrieved. If a row is compressed and the dictionary the dictionary pages have not been read when the row is encountered, the UNLOAD utility ignores this row, issues a warning message, and increments the error count. If the error count reaches the limit specified by the MAXERR option, UNLOAD terminates with an error message.

If the image copy data set is an incremental copy or a copy of pieces, the FROMCOPYDDN option can be used for a DDNAME to concatenate the data set with the corresponding full image copy that contains the dictionary. For more information, see the description of the FROMCOPYDDN option on page 480.

Interpreting field specification errors

If the UNLOAD utility detects any inconsistency relating to the field specification, including a problem in data conversion or encoding while unloading a row, DB2 issues an error message. If the MAXERR option specifies a number greater than zero, the UNLOAD utility continues processing until the total number of the records in error reaches the specified MAXERR number. DB2 issues one message for each record in error and does not unload the record. For information about specific error messages, see *DB2 Messages and Codes*.

Terminating or restarting UNLOAD

For instructions on restarting a utility job, see "Restarting an online utility" on page 46.

Terminating UNLOAD: If you terminate UNLOAD using the TERM UTILITY command during the unload phase, the output records are not erased. The output data set remains incomplete until you either delete it or restart the utility job.

Restarting UNLOAD: When the source is one or more table spaces, you can restart the UNLOAD job at the partition level or at the table space level when data is unloaded from multiple table spaces using the LIST option. When you restart a terminated UNLOAD job, processing begins with the table spaces or partitions that had not yet been completed. For a table space or partitions that were being processed at termination, UNLOAD resets the output data sets and processes those table space or partitions again.

When the source is one or more image copy data sets (FROMCOPY or FROMCOPYDDN was specified), UNLOAD always starts processing from the beginning.

Concurrency and compatibility

Individual data partitions are treated as distinct source objects. Utilities operating on different partitions of the same table space are compatible.

Claims and drains

Table 104 shows which claim classes UNLOAD drains and the restrictive states the utility sets.

Table 104. Claim classes of UNLOAD operations. Use of claims and drains; restrictive states set on the target object.

Target	UNLOAD	UNLOAD PART
Table space or physical partition of a table space with SHRLEVEL REFERENCE	DW/UTRO	DW/UTRO
Table space or physical partition of a table space with SHRLEVEL CHANGE	CR/UTRW	CR/UTRW
Image copy*	CR/UTRW	CR/UTRW

Table 104. Claim classes of UNLOAD operations (continued). Use of claims and drains; restrictive states set on the target object.

Target	UNLOAD	UNLOAD PART

Legend:
- DW: Drain the write claim class, concurrent access for SQL readers
- UTRO: Utility restrictive state, read only access allowed
- CR: Claim read, concurrent access for SQL writers and readers
- UTRW: Utility restrictive state; read/write access allowed

Note: * If the target object is an image copy, the UNLOAD utility applies CR/UTRW to the corresponding table space or physical partitions in order to prevent the table space from being dropped while data are unloaded from the image copy, even though the UNLOAD utility does not access the data in the table space.

Compatibility

The compatibilities of the UNLOAD utility and the other utilities on the same target objects are shown in Table 105. If the SHRLEVEL REFERENCE option is specified, only SQL read operations are allowed on the same target objects; otherwise SQL INSERT, DELETE, and UPDATE are allowed, in addition. If the target object is an image copy, INSERT, DELETE, and UPDATE are always allowed on the corresponding table space. In any case, DROP or ALTER can not be applied to the target object while the UNLOAD utility is running.

Table 105. Unload compatibility

Action	UNLOAD SHRLEVEL REFERENCE	UNLOAD SHRLEVEL CHANGE	FROM IMAGE COPY
CHECK DATA DELETE YES	No	No	Yes
CHECK DATA DELETE NO	Yes	Yes	Yes
CHECK INDEX	Yes	Yes	Yes
CHECK LOB	Yes	Yes	Yes
COPY INDEXSPACE	Yes	Yes	Yes
COPY TABLESPACE	Yes	Yes	Yes*
DIAGNOSE	Yes	Yes	Yes
LOAD SHRLEVEL NONE	No	No	Yes
LOAD SHRLEVEL CHANGE	No	Yes	Yes
MERGECOPY	Yes	Yes	No
MODIFY RECOVERY	Yes	Yes	No
MODIFY STATISTICS	Yes	Yes	Yes
QUIESCE	Yes	Yes	Yes
REBUILD INDEX	Yes	Yes	Yes
RECOVER (no options)	No	No	Yes
RECOVER TOCOPY or TORBA	No	No	Yes

Table 105. Unload compatibility (continued)

Action	UNLOAD SHRLEVEL REFERENCE	UNLOAD SHRLEVEL CHANGE	FROM IMAGE COPY
RECOVER ERROR RANGE	No	No	Yes
REORG INDEX	Yes	Yes	Yes
REORG TABLESPACE UNLOAD CONTINUE or PAUSE	No	No	Yes
REORG TABLESPACE UNLOAD ONLY or EXTERNAL	Yes	Yes	Yes
REPAIR DUMP or VERIFY	Yes	Yes	Yes
REPAIR LOCATE KEY or RID DELETE or REPLACE	No	No	Yes
REPAIR LOCATE TABLESPACE PAGE REPLACE	No	No	Yes
REPAIR LOCATE INDEX PAGE REPLACE	Yes	Yes	Yes
REPORT	Yes	Yes	Yes
RUNSTATS TABLESPACE	Yes	Yes	Yes
RUNSTATS INDEX	Yes	Yes	Yes
STOSPACE	Yes	Yes	Yes

Note (*): If the same data set is used as the output from the COPY utility and as the input data set of the UNLOAD utility, unexpected results can occur.

Sample control statements

Example 1: Unloading all rows and columns of a table space. This example shows the JCL for unloading rows (all columns) whose WORKDEPT column matches 'D11' and SALARY > 25000 from the EMP table in DSN8D71A.DSN8S71E.

```
//STEP1    EXEC DSNUPROC,UID='SMPLUNLD',UTPROC='',SYSTEM='V71A'
//SYSREC   DD DSN=USERID.SMPLUNLD.SYSREC,
//         DISP=(NEW,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(TRK,(2,1))
//SYSPUNCH DD DSN=USERID.SMPLUNLD.SYSPUNCH,
//         DISP=(NEW,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(TRK,(1,1))
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
  UNLOAD TABLESPACE DSN8D71A.DSN8S71E
    FROM TABLE DSN8710.EMP
    WHEN (WORKDEPT = 'D11' AND SALARY > 25000)
```

Example 2: Unloading specific columns using a field specification list. This example shows unloading specific columns in a desired order using a field specification list.

```
UNLOAD TABLESPACE DSN8D71A.DSN8S71E NOPAD
  FROM TABLE DSN8710.EMP
    (EMPNO, LASTNAME, SALARY DECIMAL EXTERNAL)
  WHEN (WORKDEPT = 'D11' AND SALARY > 25000)
```

The output from the previous example might look similar to the following output:

```
000060@@STERN#   32250.00
000150@@ADAMSON#   25280.00
000200@@BROWN#   27740.00
000220@@LUTZ#   29840.00
200220@@JOHN#   29840.00
```

In the previous output, '@@' before the last name represents the length field for the VARCHAR field LASTNAME (for example, X'0005' for STERN) and '#' is the NULL indicator byte for the nullable SALARY field. Because the SALARY column is declared as DECIMAL (9,2) on the table, the default output length of the SALARY field is 11 (9 digits + sign + decimal point) not including the NULL indicator byte. LASTNAME was unloaded as a variable length field because the NOPAD option was specified.

Example 3: UNLOAD JCL to unload data from a partitioned table space in parallel. If the source table space is partitioned and each partition is associated with an output data set using a TEMPLATE utility statement, the UNLOAD job will run in parallel in a multi-processor environment. The example below shows a case where data is unloaded from all partitions to data sets that will be dynamically allocated. The number of parallel tasks will be determined by the number of available processors.

Assume that table space TDB1.TSP1 has three partitions and contains table TCRT.TTBL.

```
//STEP1    EXEC DSNUPROC,UID='SMPLUNLD',UTPROC='',SYSTEM='V71A'
//SYSPUNCH DD DSN=USERID.SMPLUNLD.SYSPUNCH,
//         DISP=(NEW,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(TRK,(1,1))
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
  TEMPLATE UNLDDS   DSN &USERID..SMPLUNLD.&TS..P&PART.
           UNIT SYSDA  DISP (NEW,CATLG,CATLG) SPACE (2,1) CYL
  UNLOAD TABLESPACE TDB1.TSP1
    UNLDDN UNLDDS
    FROM TABLE TCRT.TTBL
```

As a result, assuming that the user ID is USERID, the job creates three data sets similar to the following output:

```
   USERID.SMPLUNLD.TSP1.P00001  ... contains rows from partition 1
   USERID.SMPLUNLD.TSP1.P00002  ... contains rows from partition 2
   USERID.SMPLUNLD.TSP1.P00003  ... contains rows from partition 3
```

Example 4: Selecting UNLOAD partitions using a LISTDEF utility statement. The following example shows unloading data from partitions one and three using a LISTDEF specification.

```
//SAMPJOB  JOB ...
//STEP1    EXEC DSNUPROC,UID='SMPLUNLD',UTPROC='',SYSTEM='V71A'
//SYSPUNCH DD DSN=USERID.SMPLUNLD.SYSPUNCH,
//         DISP=(NEW,CATLG,CATLG),
//         UNIT=SYSDA,SPACE=(TRK,(1,1))
```

```
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
  LISTDEF  UNLDLIST
           INCLUDE  TABLESPACE TDB1.TSP1 PARTLEVEL(1)
           INCLUDE  TABLESPACE TDB1.TSP1 PARTLEVEL(3)
  TEMPLATE UNLDDS   DSN &USERID..SMPLUNLD.&TS..P&PART.
           UNIT SYSDA  DISP (NEW,CATLG,CATLG) SPACE (2,1) CYL
  UNLOAD LIST UNLDLIST              -- LIST name
    UNLDDN UNLDDS                   -- TEMPLATE name
```

The result will be stored in the following two data sets:

- USERID.SMPLUNLD.TSP1.P00001 ... contains rows from partition 1
- USERID.SMPLUNLD.TSP1.P00003 ... contains rows from partition 3

Example 5: Unloading multiple table spaces using LISTDEF. The following example illustrates unloading data from multiple table spaces using a single UNLOAD statement. Assume that if the database TDB1 contains two table spaces that can be expressed by the pattern-matching string 'TSP*', (for example, TSP1 and TSP2), then two output data sets will be created.

```
//SAMPJOB  JOB ...
//STEP1    EXEC DSNUPROC,UID='SMPLUNLD',UTPROC='',SYSTEM='V71A'
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
  LISTDEF  UNLDLIST
           INCLUDE  TABLESPACE TDB1.TSP*
  TEMPLATE UNLDDS   DSN &USERID..SMPLUNLD.&TS.
           UNIT SYSDA  DISP (NEW,CATLG,CATLG) SPACE (2,1) CYL
  TEMPLATE PUNCHDS  DSN &USERID..SMPLPUNC.&TS.
           UNIT SYSDA  DISP (NEW,CATLG,CATLG) SPACE (1,1) CYL
  UNLOAD LIST UNLDLIST
    PUNCHDDN PUNCHDS              -- TEMPLATE name
    UNLDDN   UNLDDS               -- TEMPLATE name
```

This example will create two output data sets:

- USERID.SMPLUNLD.TSP1 ... rows from table space TDB1.TSP1
- USERID.SMPLUNLD.TSP2 ... rows from table space TDB1.TSP2

Part 3. Stand-alone utilities

Chapter 30. Invoking stand-alone utilities . . . 527
Creating utility statements and EXEC PARM parameters . . . 527
Creating utility control statements . . . 527
Control statement coding rules . . . 527
Specifying options using the JCL EXEC PARM parameter . . . 527
Following OS/390 JCL EXEC PARM specification rules . . . 527
Example of an option description . . . 528

Chapter 31. DSNJLOGF (Preformat Active Log) . . . 529
Before running DSNJLOGF . . . 529
Environment . . . 529
Control statement . . . 529
Sample control statement . . . 529
DSNJLOGF output . . . 529

Chapter 32. DSNJU003 (Change Log Inventory) . . . 531
Syntax and options of the control statement . . . 531
DSNJU003 (change log inventory) syntax diagram . . . 531
Option descriptions . . . 533
Before running DSNJU003 . . . 539
Environment . . . 539
Authorization required . . . 540
Control statement . . . 540
Required and optional data sets . . . 540
Optional statements . . . 540
Using DSNJU003 . . . 541
Running DSNJU003 . . . 541
Making changes for active logs . . . 541
Making changes for archive logs . . . 543
Creating a conditional restart control record . . . 543
Deleting log data sets with errors . . . 543
Altering references to NEWLOG and DELETE data sets . . . 544
Specifying the NEWCAT statement . . . 544
Renaming DB2 system data sets . . . 545
Renaming DB2 active log data sets . . . 545
Renaming DB2 archive log data sets . . . 546
Sample control statements . . . 546

Chapter 33. DSNJU004 (Print Log Map) . . . 547
Syntax and options of the control statement . . . 547
DSNJU004 (print log map) syntax diagram . . . 547
Option descriptions . . . 547
Before running DSNJU004 . . . 548
Environment . . . 548
Authorization required . . . 548
Control statement . . . 548
Required and optional data sets . . . 548
Running the DSNJU004 utility . . . 549
Recommendations . . . 549
Sample control statement . . . 549
DSNJU004 (Print Log Map) output . . . 549
Timestamps in the BSDS . . . 550
Active log data set status . . . 551

© Copyright IBM Corp. 1983, 2001

Reading conditional restart control records . . . 554

Chapter 34. DSN1CHKR . . . 555
Syntax and options of the control statement . . . 555
DSN1CHKR syntax diagram . . . 555
Option descriptions . . . 555
Before running DSN1CHKR . . . 557
Environment . . . 557
Authorization required . . . 557
Control statement . . . 557
Restrictions . . . 557
Running DSN1COPY before DSN1CHKR . . . 557
Running DSN1CHKR on a valid table space . . . 558
Sample control statements . . . 558
DSN1CHKR output . . . 561

Chapter 35. DSN1COMP . . . 563
Syntax and options of the control statement . . . 563
DSN1COMP syntax diagram . . . 563
Option descriptions . . . 563
Before running DSN1COMP . . . 566
Environment . . . 566
Authorization required . . . 566
Control statement . . . 566
Recommendation . . . 567
Using DSN1COMP . . . 567
Estimating compression savings achieved by REORG . . . 567
Including free space in compression calculations . . . 568
Running DSN1COMP on a table space with identical data . . . 568
Sample control statements . . . 568
DSN1COMP output . . . 568
Message DSN1941 . . . 568
Sample DSN1COMP report . . . 569

Chapter 36. DSN1COPY . . . 571
Syntax and options of the control statement . . . 571
DSN1COPY syntax diagram . . . 572
Option descriptions . . . 572
Before running DSN1COPY . . . 576
Environment . . . 576
Authorization required . . . 576
Control statement . . . 577
Required data sets . . . 577
Defining the input data set . . . 579
Defining the output data set . . . 580
Naming the output data set . . . 580
Adding additional volumes for SYSUT2 . . . 581
Restrictions . . . 581
Recommendations . . . 581
Printing with DSN1PRNT instead of DSN1COPY . . . 581
Determining page size and DSSIZE . . . 581
Using DSN1COPY . . . 582
Altering a table before running DSN1COPY . . . 582
Checking for inconsistent data . . . 582
Translating DB2 internal identifiers . . . 582
Using an image copy as input to DSN1COPY . . . 583

Resetting page log RBAs . . . 583
Copying multiple data set table spaces . . . 583
Restoring indexes with DSN1COPY . . . 583
Restoring table spaces with DSN1COPY . . . 584
Printing with DSN1COPY . . . 584
Copying tables from one subsystem to another . . . 584
Sample control statements . . . 585
DSN1COPY output . . . 586

Chapter 37. DSN1LOGP . . . 587
Syntax and options of the control statement . . . 587
DSN1LOGP syntax diagram . . . 588
Option descriptions . . . 588
Before running DSN1LOGP . . . 593
Environment . . . 593
Authorization required . . . 594
Control statement . . . 594
Required data sets . . . 594
Identifying log data sets . . . 594
Using DSN1LOGP . . . 595
Reading archive log data sets on tape . . . 595
Locating table and index identifiers . . . 596
Sample control statements . . . 596
DSN1LOGP output . . . 599
Reviewing DSN1LOGP output . . . 599
Description of the summary report . . . 599
Description of the detail report . . . 602
Interpreting data propagation information in the summary report . . . 606
Interpreting error codes . . . 607

Chapter 38. DSN1PRNT . . . 609
Syntax and options of the control statement . . . 609
DSN1PRNT syntax diagram . . . 610
Option descriptions . . . 610
Before running DSN1PRNT . . . 614
Environment . . . 614
Authorization required . . . 614
Control statement . . . 614
Recommendations . . . 615
Printing with DSN1PRNT instead of DSN1COPY . . . 615
Determining page size and DSSIZE . . . 615
Sample control statements . . . 615
DSN1PRNT output . . . 616

Chapter 39. DSN1SDMP . . . 617
Syntax and options of the control statement . . . 617
DSN1SDMP syntax diagram . . . 617
Option descriptions . . . 617
Before running DSN1SDMP . . . 620
Environment . . . 620
Authorization required . . . 620
Control statement . . . 621
Using DSN1SDMP . . . 621
Assigning buffers . . . 621
Generating a dump . . . 622
Stopping or modifying DSN1SDMP traces . . . 622

Sample control statements 623
DSN1SDMP output 624

Chapter 30. Invoking stand-alone utilities

This chapter contains procedures and guidelines for creating utility control statements and EXEC PARM parameters for invoking the stand-alone utilities.

Creating utility statements and EXEC PARM parameters

Utility control statements and parameters define the function a utility job performs. Some stand-alone utilities read the control statements from an input stream, and others obtain the function definitions from JCL EXEC PARM parameters.

Creating utility control statements

You can create the utility control statements with the ISPF/PDF edit function. After you create the control statements, save them in a sequential or partitioned data set.

The following utilities read control statements from the input stream file of the specified DD name:

Utility	DD name
DSNJU003 (Change Log Inventory)	SYSIN
DSNJU004 (Print Log Map)	SYSIN (optional)
DSN1LOGP	SYSIN
DSN1SDMP	SDMPIN

Control statement coding rules

Utility control statements are read from the DD name input stream. The statements in that stream must conform to these rules:

- The logical record length (LRECL) must be 80 characters. Columns 73 through 80 are ignored.
- The records are concatenated into a single stream before being parsed. No concatenation character is necessary.
- The SYSIN stream can contain multiple utility control statements.

Specifying options using the JCL EXEC PARM parameter

The following stand-alone utilities obtain function options from the EXEC PARM parameter:

DSN1CHKR
DSN1COMP
DSN1COPY
DSN1PRNT

Following OS/390 JCL EXEC PARM specification rules

The parameters you specify must obey these OS/390 JCL EXEC PARM parameter specification rules:

- Enclose multiple subparameters in single quotes or parentheses, separating subparameters with commas. For example:

```
//name   EXEC PARM='ABC,...,XYZ'
```

- The total length cannot exceed 100 characters.
- Blanks are not allowed within the parameter specification.

© Copyright IBM Corp. 1983, 2001

To specify the parameter across multiple lines:

1. Enclose it in parentheses
2. End the first line with a subparameter, followed by a comma
3. Continue the subparameters on the next line, beginning before column 17.

For example:

```
//stepname  EXEC PARM=(ABC,...LMN,
           OPQ,...,XYZ)
```

Example of an option description

Where the syntax of each utility control statement is described, parameters are indented under the option keyword they must follow. Here is an example:

AFTER*(integer)*

Specifies that the action is to be performed after the trace point is reached *integer* times.

integer must be between 1 and 32767. The **default** is **(1)**.

In the example, AFTER is an option keyword, and *integer* is a parameter. Values of parameters are usually enclosed in parentheses. The syntax diagrams for utility control statements show parentheses where they are required.

Chapter 31. DSNJLOGF (Preformat Active Log)

When writing to an active log data set for the first time, DB2 must preformat a VSAM control area before writing the log records. The DSNJLOGF utility avoids this delay by preformatting the active log data sets before bringing them online to DB2.

Before running DSNJLOGF

This section contains information you need to be aware of prior to running DSNJLOGF.

Environment

Run DSNJLOGF as an MVS job.

Control statement

See "Sample control statement" for an example of using DSNJLOGF to preformat the active log data sets.

Required data sets: DSNJLOGF recognizes DD statements with the following DD names.

SYSUT1 Defines the newly defined active log data set to be preformatted. The data set must be an empty VSAM linear data set.

SYSPRINT Defines the print spool class or data set for print output. The logical record length (LRECL) is 132.

Sample control statement

The following sample control statement preformats the active log data sets.

```
//JOBLIB DD  DSN=DSN710.SDSNLOAD,DISP=SHR
//STEP1    EXEC PGM=DSNJLOGF
//SYSPRINT DD SYSOUT=A
//SYSUDUMP DD SYSOUT=A
//SYSUT1   DD DSN=DSNC710.LOGCOPY1.DS01,DISP=SHR
//STEP2    EXEC PGM=DSNJLOGF
//SYSPRINT DD SYSOUT=A
//SYSUDUMP DD SYSOUT=A
//SYSUT1   DD DSN=DSNC710.LOGCOPY1.DS02,DISP=SHR
//STEP3    EXEC PGM=DSNJLOGF
//SYSPRINT DD SYSOUT=A
//SYSUDUMP DD SYSOUT=A
//SYSUT1   DD DSN=DSNC710.LOGCOPY2.DS01,DISP=SHR
//STEP4    EXEC PGM=DSNJLOGF
//SYSPRINT DD SYSOUT=A
//SYSUDUMP DD SYSOUT=A
//SYSUT1   DD DSN=DSNC710.LOGCOPY2.DS02,DISP=SHR
```

DSNJLOGF output

```
DSNJ991I DSNJLOGF START OF LOG DATASET PREFORMAT FOR JOB LOGFRMT  STEP1
DSNJ992I DSNJLOGF LOG DATA SET NAME = DSNC710.LOGCOPY1.DS01
DSNJ996I DSNJLOGF LOG PREFORMAT COMPLETED SUCCESSFULLY, 00015000
         RECORDS FORMATTED
```

© Copyright IBM Corp. 1983, 2001

Chapter 32. DSNJU003 (Change Log Inventory)

The DSNJU003 stand-alone utility changes the bootstrap data sets (BSDSs). You can use the utility to:

- Add or delete active or archive log data sets
- Add or delete checkpoint records
- Create a conditional restart control record to control the next start of the DB2 subsystem
- Change the VSAM catalog name entry in the BSDS
- Modify the communication record in the BSDS
- Modify the value for the highest-written log RBA value (relative byte address within the log) or the highest-offloaded RBA value

Syntax and options of the control statement

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

DSNJU003 (change log inventory) syntax diagram

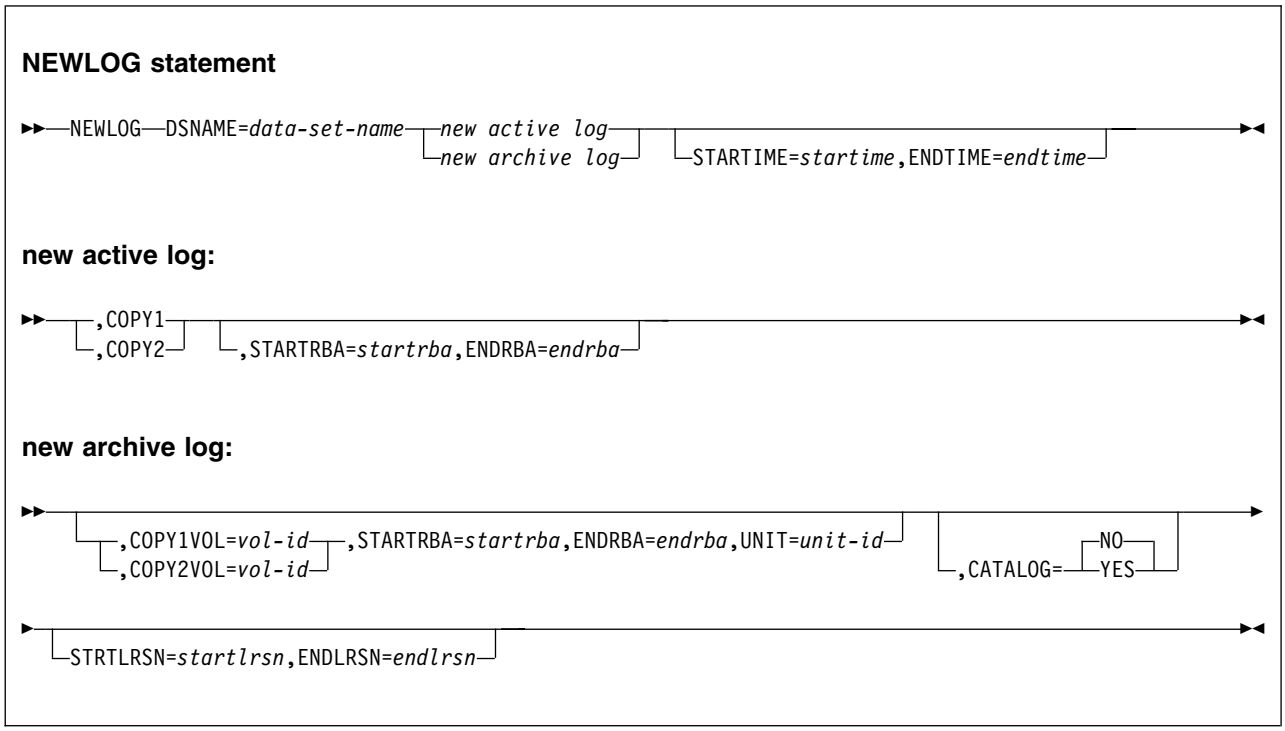

© Copyright IBM Corp. 1983, 2001

DSNJU003 (Change Log Inventory)

DELETE statement

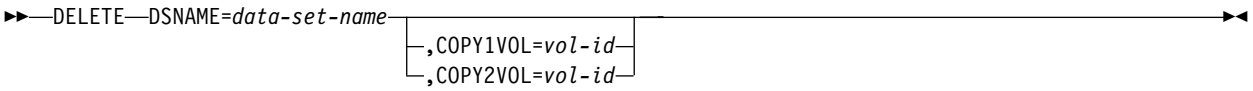

CRESTART statement

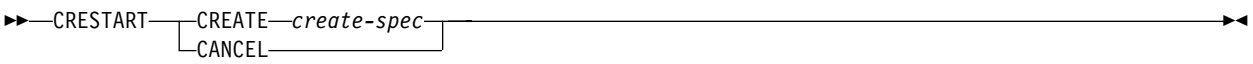

create-spec:

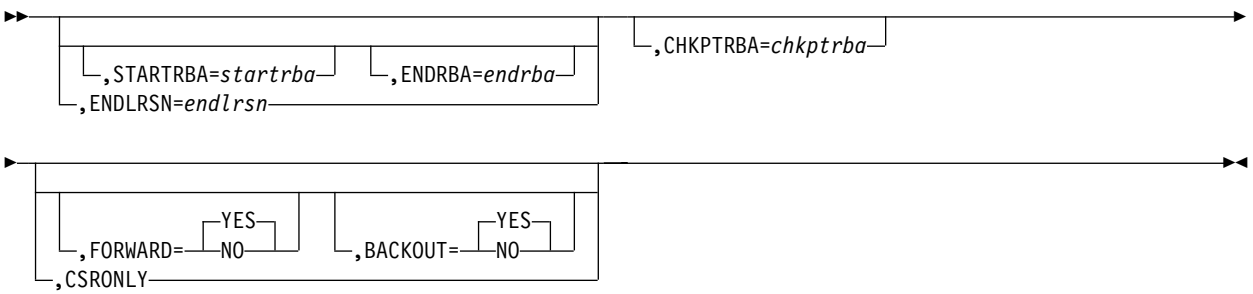

NEWCAT statement

DDF statement

CHECKPT statement

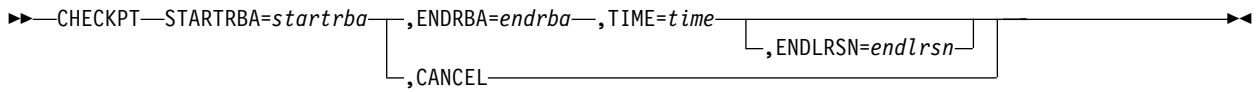

HIGHRBA statement

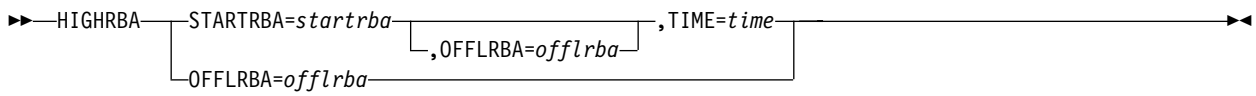

Option descriptions

For a description of how utility statements are parsed, and how to read a list of option identifiers and specifications like the one that follows, see “Control statement coding rules” on page 527.

NEWLOG Declares one of the following data sets:

- A VSAM data set that is available for use as an active log data set.

  Use only the keywords DSNAME=, COPY1, and COPY2.
- An active log data set that is replacing one that encountered an I/O error.

  Use only the keywords DSNAME=, COPY1, COPY2, STARTRBA=, and ENDRBA=.
- An archive log data set volume.

  Use only the keywords DSNAME= ,COPY1VOL=, COPY2VOL=, STARTRBA=, ENDRBA=, UNIT=, CATALOG=, STRTLRSN=, and ENDLRSN=.

  If you create an archive log data set and add it to the BSDS with this utility, you can specify a name that DB2 might also generate. DB2 generates archive log data set names of the form DSNCAT.ARCHLOG*x*.A*nnnnnnn* where:
  - DSNCAT and ARCHLOG are parts of the data set prefix you specified on install panels DSNTIPA2 and DSNTIPH.
  - *x* is 1 for the first copy of the logs and 2 for the second copy.
  - A*nnnnnnn* represents the series of low-level qualifiers DB2 generates for archive log data set names, beginning with A0000001, and incrementing to A0000002, A0000003, and so forth.

  For data sharing, the naming convention is DSNCAT.ARCHLOG1 or DSNCAT.DSN1.ARCLG1.

  If you do specify a name using the same naming convention as DB2, you receive a dynamic allocation error when DB2 generates that name. The error message, DSNJ103I, is issued

once. DB2 then increments the low-level qualifier to generate the next data set name in the series and offloads to it the next time DB2 archives. (The active log that previously was not offloaded is offloaded to this data set.)

The newly-defined active logs cannot specify a start and end LRSN. The start and end LRSN for new active logs that contain active log data are read at DB2 start-up time from the new active log data sets specified in the Change Log Inventory NEWLOG statements. For new archive logs defined with Change Log Inventory, the user must specify the start and end RBAs. For data sharing, the user must also specify the start and end LRSNs. DB2 startup does not attempt to find these values from the new archive log data sets.

DELETE Deletes all information about the specified log data set or data set volume from the bootstrap data sets.

CRESTART Controls the next restart of DB2, either by creating a new conditional restart control record or by canceling the one that is currently active.

Attention: This statement can override DB2's efforts to maintain data in a consistent state. Do not use this statement without understanding the conditional restart process, which is described in Part 4 (Volume 1) of *DB2 Administration Guide*.

NEWCAT Changes the VSAM catalog name in the BSDS.

DDF Updates the LOCATION, LUNAME, and PASSWORD values in the BSDS. If you use this statement to insert new values into the BSDS, you must include at least the LOCATION and LUNAME in the DDF statement. To update an existing set of values, you need only include those values you want to change. The DDF record cannot be deleted from the BSDS after it has been added, it can only be modified.

NOPASSWD removes the DDF password from the DDF record in the BSDS. No other keywords can be used with NOPASSWD.

CHECKPT Allows updating of the checkpoint queue with the start checkpoint and end checkpoint log records.

Attention: This statement can override DB2's efforts to maintain data in a consistent state. Do not use the statement without understanding the conditional restart and checkpoint processing processes, which are described in Part 4 (Volume 1) of *DB2 Administration Guide*.

HIGHRBA Updates the highest-written log RBA in either the active or archive log data sets.

Attention: This statement can override DB2's efforts to maintain data in a consistent state. Do not use the statement without understanding the conditional restart process, which is described in Part 4 (Volume 1) of *DB2 Administration Guide*.

DSNAME=*data-set-name*
Specifies a log data set.

data-set-name can be up to 44 characters long.

STARTIME=*startime*

Enables you to record the start time of the RBA in the BSDS. This is an optional field. The timestamp format with valid values in parentheses is as follows:

yyyydddhhmmsst

where:

yyyy	Indicates the year (1989-2099).
ddd	Indicates the day of the year (0-365; 366 in leap years).
hh	Indicates the hour (0-23).
mm	Indicates the minutes (0-59).
ss	Indicates the seconds (0-59).
t	Indicates tenths of a second.

If fewer than 14 digits are specified for the STARTIME or ENDTIME parameter, trailing zeros are added.

If STARTIME is specified, the ENDTIME, STARTRBA and ENDRBA must also be specified.

ENDTIME=*endtime*

Enables you to record the end time of the RBA in the BSDS. This is an optional field. For the timestamp format, see the STARTIME option. The ENDTIME value must be greater than or equal to the value of STARTIME.

COPY1 Makes the data set an active log copy-1 data set.

COPY2 Makes the data set an active log copy-2 data set.

STARTRBA=*startrba*

startrba is a hexadecimal number of up to 12 characters. If you use fewer than 12 characters, leading zeros are added. *startrba* must end with '000' or DB2 returns a DSNJ4381 error message. The RBA can be obtained from messages or by printing the log map.

On the NEWLOG statement, *startrba* gives the log RBA of the beginning of the replacement active log data set or the archive log data set volume specified by DSNAME.

On the CRESTART statement, *startrba* is the earliest RBA of the log to be used during restart. If you omit STARTRBA, DB2 determines the beginning of the log range.

On the CHECKPT statement, *startrba* indicates the start checkpoint log record.

STARTRBA is required when STARTIME is specified.

for the *startrba* format, see the NEWLOG statement.

On the HIGHRBA statement, *startrba* denotes the log RBA of the highest-written log record in the active log data sets.

For the *startrba* format, see the NEWLOG statement.

ENDRBA=*endrba*

endrba is a hexadecimal number of up to 12 characters. If you use fewer than 12 characters, leading zeros are added. *endrba* must end with 'FFF' or DB2 will return a DSNJ4381 error message.

On the NEWLOG statement, *endrba* gives the log RBA (relative byte address within the log) of the end of the replacement active log data set or the archive log data set volume specified by DSNAME.

On the CRESTART statement, *endrba* is the last RBA of the log that is to be used during restart, and it is also the starting RBA of the next active log written after restart. Any log information in the bootstrap data set and the active logs with an RBA greater than *endrba* is discarded. If you omit ENDRBA, DB2 determines the end of the log range.

The value of ENDRBA must be a multiple of 4096. (The hexadecimal value must end in 000.) Also, the value must be greater than or equal to the value of STARTRBA. If STARTRBA and ENDRBA are equal, the next restart is a cold start; that is, no log records are processed during restart. The specified RBA becomes the beginning RBA of the new log.

On the CHECKPT statement , *endrba* indicates the end checkpoint log record that corresponds to the start checkpoint log record.

For the *endrba* format, see the NEWLOG statement.

COPY1VOL=*vol-id*

vol-id is the volume serial of the copy-1 archive log data set that is specified after DSNAME.

COPY2VOL=*vol-id*

vol-id is the volume serial of the copy-2 archive log data set that is specified after DSNAME.

UNIT=*unit-id* *unit-id* is the device type of the archive log data set that is named after DSNAME.

CATALOG Indicates whether the archive log data set is cataloged.

NO Indicates that the archive log data set is not cataloged. All subsequent allocations of the data set are made using the unit and volume information specified on the statement.

YES Indicates that the archive log data set is cataloged. All subsequent allocations of the data set are made using the catalog.

DB2 requires that all archive log data sets on DASD be cataloged. Select CATALOG=YES if the archive log data set is on DASD.

STRTLRSN=*startlrsn*

On the NEWLOG statement, *startlrsn* is the LRSN in log record header of the first complete log record on the new archive data set. *startlrsn* is a hexadecimal number of up to 12 characters. If you use fewer than 12 characters, leading zeros are added. In a data sharing environment, run the Print Log Map utility to find an archive log data set and start and end RBAs and LRSNs.

ENDLRSN=*endlrsn*

endlrsn is a hexadecimal number of up to 12 characters. If you use fewer than 12 characters, leading zeros are added. In a data

sharing environment, run the Print Log Map utility to find an archive log data set and start and end RBAs and LRSNs.

On the NEWLOG statement, *endlrsn* is the LRSN in log record header of the last log record on the new archive data set.

On the CRESTART statement, *endlrsn* is the LRSN of the last log record to be used during restart. Any log information in the bootstrap data set and the active logs with an LRSN greater than *endlrsn* is discarded. If you omit ENDLRSN, DB2 determines the end of the log range.

The ENDLRSN option is valid only in a data sharing environment. It cannot be specified with STARTRBA or ENDRBA.

On the CHECKPT statement, *endlrsn* is the LRSN of the end checkpoint log record.

CREATE Creates a new conditional restart control record. When the new record is created, the previous control record becomes inactive.

CANCEL **On the CRESTART statement,** CANCEL deactivates the currently active conditional restart control record. The record remains in the BSDS as historical information.

No other keyword can be used with CANCEL.

On the CHECKPT statement, CANCEL deletes the checkpoint queue entry that contains a starting RBA that matches the parameter specified by the STARTRBA keyword.

CHKPTRBA=*chkptrba*
Is the log RBA of the start of the checkpoint record that is to be used during restart.

If you use STARTRBA or ENDRBA, and you do not use CHKPTRBA, the DSNJU003 utility selects the RBA of an appropriate checkpoint record. If you do use CHKPTRBA, you override the value selected by the utility. However, *chkptrba* must be in the range determined by *startrba* and *endrba* or their default values. If possible, do not use CHKPTRBA; let the utility determine the RBA of the checkpoint record.

CHKPTRBA=0 overrides any selection by the utility; at restart, DB2 attempts to use the most recent checkpoint record taken.

FORWARD= Indicates whether to use the forward-log-recovery phase of DB2 restart, which reads the log forward to recover any units of recovery that were in one of the following two states when DB2 was last stopped:

- Indoubt (the units of recovery had finished the first phase of commit, but had not started the second phase)
- In-commit (had started but had not finished the second phase of commit)

YES Allows forward-log recovery.

If you specify a cold start (by using the same value for STARTRBA and ENDRBA), no recovery processing is performed.

NO Terminates forward-log recovery before log records are processed.

BACKOUT= Indicates whether to use the backward-log-recovery phase of DB2 restart, which rolls back any units of recovery that were in one of the following two states when DB2 was last stopped:

- Inflight (did not complete the first phase of commit)
- In-abort (had started but not finished an abort)

YES Allows backward-log recovery.

If you specify a cold start (by using the same value for STARTRBA and ENDRBA), no recovery processing is performed.

NO Terminates backward-log recovery before log records are processed.

CSRONLY Performs only the first and second phases of restart processing (log initialization and current-status rebuild). After these phases, the system status is displayed and restart terminates. Some parts of the log initialization are not performed, including any updating of the log and display of STARTRBA and ENDRBA information.

When DB2 is restarted with this option in effect, the conditional restart control record is not deactivated. To prevent the control record from remaining active, use the DSNJU003 utility again with CRESTART CANCEL, or with CRESTART CREATE to create a new active control record.

VSAMCAT=*catalog-name*
Changes the VSAM catalog name entry in the BSDS.

catalog-name can be up to eight characters long. The first character must be alphabetic, and the remaining characters can be alphanumeric.

LOCATION=*location-name*
Changes the LOCATION value in the BSDS.

location-name specifies the name of your local DB2 site.

LUNAME=*luname*
Changes the LUNAME value in the BSDS.

The LUNAME in the BSDS must always contain the value that identifies your local DB2 subsystem to the VTAM® network.

PASSWORD= The DDF password follows VTAM convention, but DB2 restricts it to one to eight alphanumeric characters. The first character must be either a capital letter or an alphabetic extender. The remaining characters can consist of alphanumeric characters and alphabetic extenders.

password Specifying a *password* is optional. It assigns a password to the distributed data facility communication record that establishes communications for a distributed data environment. See *VTAM for MVS/ESA Resource Definition Reference* for a description of the PRTCT=*password* option on the APPL definition statement that is used to define DB2 to VTAM.

NOPASSWD Removes the archive password protection for all archives created after this operation. It also

removes a previously existing password from the DDF record. No other keyword can be used with NOPASSWD.

GENERIC=*gluname*
Replaces the value of the DB2 GENERIC LUNAME subsystem parameter in the BSDS.

NGENERIC Changes the DB2 GENERIC LUNAME to binary zeros in the BSDS, indicating that no VTAM generic LU name support is requested.

PORT Identifies the TCP/IP port number used by DDF to accept incoming connection requests. This value must be a decimal number between 0 and 65534; zero indicates that DDF's TCP/IP support is to be deactivated.

If DB2 is part of a data sharing group, all the members of the DB2 data sharing group must have the same value for PORT.

RESPORT Identifies the TCP/IP port number used by DDF to accept incoming DRDA® two-phase commit resynchronization requests. This value must be a decimal number between 0 and 65534; zero indicates that DDF's TCP/IP support is to be deactivated. If RESPORT is non-zero, RESPORT must not be the same as the value supplied on PORT.

For data sharing DB2 systems, RESPORT must be uniquely assigned to each DB2 member, so that no two DB2 members use the same TCP/IP port for two-phase commit resynchronization.

TIME=*time* **On the CHECKPT statement**, TIME gives the time the start checkpoint record was written.

For timestamp format, see the STARTIME option on the NEWLOG statement on page 535.

On the HIGHRBA statement, TIME specifies when the log record with the highest RBA was written to the log.

For timestamp format, see the STARTIME option on the NEWLOG statement on page 535.

OFFLRBA=*offlrba*
Specifies the highest offloaded RBA in the archive log.

offlrba is a hexadecimal number of up to 12 characters. If you use fewer than 12 characters, leading zeros are added. The value must end with hexadecimal 'FFF'.

Before running DSNJU003

This section contains information you need to be aware of prior to running DSNJU003.

Environment

The utility should be executed only as a batch job when DB2 is not running. It can be executed when DB2 is running, but results can be inconsistent. Changing a BSDS for a data-sharing member using DSNJU003 might cause a log read request from another data-sharing member to fail. The failure occurs only if the second member tries to access the changed BSDS before the first member is started.

Authorization required

The authorization ID of the DSNJU003 job must have the requisite RACF authorization.

Control statement

See “Syntax and options of the control statement” on page 531 for DSNJU003 syntax and option descriptions.

Required and optional data sets

DSNJU003 recognizes DD statements with the following ddnames:

JOBCAT

STEPCAT Specifies the catalog in which the bootstrap data sets (BSDSs) are cataloged. This statement is optional. Typically, the high-level qualifier of the BSDS name points to the integrated facility catalog that contains an entry for the BSDS.

SYSUT1 Specifies and allocates the bootstrap data set. This statement is required.

SYSUT2 Specifies and allocates a second copy of the bootstrap data set. This statement is required if you use dual BSDSs.

Dual BSDSs and DSNJU003: With each execution of DSNJU003, the BSDS timestamp field is updated with the current system time. If you run DSNJU003 separately for each copy of a dual copy BSDS, the timestamp fields are not synchronized, and DB2 fails at startup. If you changed the contents of the BSDS copy by running DSNJU003, DB2 issues error message DSNJ122I. Therefore, if you use DSNJU003 to update dual copy BSDSs, update both BSDSs within a single execution of DSNJU003.

SYSPRINT Specifies a data set for print output. This statement is required. The logical record length (LRECL) is 125.

SYSIN Specifies the input data set for statements. This statement is required. The logical record length (LRECL) is 80.

Optional statements

The Change Log Inventory utility provides the following statements:

- NEWLOG
- DELETE
- SYSTEMDB
- CRESTART
- NEWCAT
- DDF
- CHECKPT
- HIGHRBA

You can specify any statement one or more times. In each statement, separate the operation name from the first parameter by one or more blanks. You can use parameters in any order; separate them by commas with no blanks. Do not split a parameter description across two SYSIN records.

A statement containing an asterisk in column 1 is considered a comment and is ignored. However, it appears in the output listing. To include a comment or sequence number in a SYSIN record, separate it from the last comma by a blank. When a blank is encountered following a comma, the rest of the record is ignored.

During execution of DSNJU003, a significant error in any statement causes that statement and all subsequent statements to be skipped. However, all remaining statements are checked for syntax errors. Therefore, BSDS updates are not made for any operation specified in the statement in error and in any subsequent statements.

Using DSNJU003

This section describes the following tasks associated with running the DSNJU003 utility:

"Running DSNJU003"
"Making changes for active logs"
"Making changes for archive logs" on page 543
"Creating a conditional restart control record" on page 543
"Deleting log data sets with errors" on page 543
"Altering references to NEWLOG and DELETE data sets" on page 544
"Specifying the NEWCAT statement" on page 544
"Renaming DB2 system data sets" on page 545
"Renaming DB2 active log data sets" on page 545
"Renaming DB2 archive log data sets" on page 546

Running DSNJU003

The following statement executes the utility and can be included only in a batch job:

```
//EXEC   PGM=DSNJU003
```

Making changes for active logs

Adding: If an active log is in stopped status, it is not reused for output logging; however, it continues to be used for reading. To add a new active log:

1. Use the access method services DEFINE command to define new active log data sets.
2. Use DSNJLOGF to preformat the new active log data sets.
3. Use DSNJU003 to register the new data sets in the BSDS.

For example, specify:

```
NEWLOG DSNAME=DSNC710.LOGCOPY1.DS04,COPY1
NEWLOG DSNAME=DSNC710.LOGCOPY2.DS04,COPY2
```

If you are copying the contents of an old active log data set to the new one, you can also give the RBA range and the starting and ending timestamp on the NEWLOG statement.

If you are archiving to DASD and the size of your active logs has been increased, you might find it necessary to increase the size of your archive log data sets.

Deleting: To delete information about an active log data set from the BSDS, you might specify:

```
DELETE DSNAME=DSNC710.LOGCOPY1.DS01
DELETE DSNAME=DSNC710.LOGCOPY2.DS01
```

Recording: To record information about an existing active log data set in the BSDS, you might specify:

```
NEWLOG DSNAME=DSNC710.LOGCOPY2.DS05,COPY2,STARTIME=19910212205198,
             ENDTIME=19910412205200,STARTRBA=43F8000,ENDRBA=65F3FFF
```

You can insert a record of that information into the BSDS for any of these reasons:

- The data set has been deleted and is needed again.
- You are copying the contents of one active log data set to another data set (copy 1 to copy 2).
- You are recovering the BSDS from a backup copy.

Enlarging: When DB2 is inactive (down), use one of the following procedures.

If you can use the access method services REPRO command, follow these steps:

1. Stop DB2. This step is required because DB2 allocates all active log data sets when it is up.
2. Use the access method services ALTER command with the NEWNAME option to rename your active log data sets.
3. Use the access method services DEFINE command to define larger active log data sets. Refer to installation job DSNTIJIN to see the definitions that create the original active log data sets. See *DB2 Installation Guide*.

   By reusing the old data set names, you don't have to run the Change Log Inventory utility to establish new names in the BSDSs. The old data set names and the correct RBA ranges are already in the BSDSs.
4. Use the access method services REPRO command to copy the old (renamed) data sets into their respective new data sets.
5. Start DB2.

If you cannot use the access method services REPRO command, follow this procedure:

1. Ensure that all active log data sets except the current active log data sets have been archived. Active log data sets that have been archived are marked REUSABLE in Print Log Map utility (DSNJU004) output.
2. Stop DB2.
3. Rename or delete the reusable active logs. Allocate new, larger active log data sets with the same names as the old active log data sets.
4. Run the DSNJLOGF utility to preformat the new log data sets.
5. Run the Change Log Inventory utility (DSNJU003) with the DELETE statement to delete all active logs except the current active logs from the BSDS.
6. Run the Change Log Inventory utility with the NEWLOG statement to add to the BSDS the active logs that you just deleted. So that the logs are added as empty, do not specify an RBA range.
7. Start DB2.
8. Execute the ARCHIVE LOG command to cause DB2 to truncate the current active logs and switch to one of the new sets of active logs.
9. Repeat steps 2 through 7 to enlarge the active logs that were just archived.

Although it is not necessary for all log data sets to be the same size, from an operational standpoint it is more consistent and efficient. If the log data sets are not the same size, it is more difficult to track your system's logs. Space can be wasted if you are using dual data sets of different sizes because they will fill only to the size of the smallest, not using the remaining space on the larger one.

If you are archiving to disk and the size of your active logs has been increased, you might find it necessary to increase the size of your archive log data sets. You must specify less than 64 KB tracks for the primary space quantity, because of DFSMS

Direct Access Device Space Management limits. Refer to the PRIMARY QUANTITY and SECONDARY QTY fields on installation panel DSNTIPA to modify the primary and secondary allocation space quantities. See *DB2 Installation Guide* for more information.

Making changes for archive logs

Adding: When the recovery of an object depends on reading an existing archive log data set, the BSDS must contain information about that data set, so that the recovery job can find it. To register information about an existing archive log data set in the BSDS, you might specify:

```
NEWLOG DSNAME=DSNC710.ARCHLOG1.D89021.T2205197.A0000015,COPY1VOL=DSNV04,
UNIT=TAPE,STARTRBA=3A190000,ENDRBA=3A1F0FFF,CATALOG=NO
```

Deleting: To delete an entire archive log data set from one or more volumes, you might specify:

```
DELETE DSNAME=DSNC710.ARCHLOG1.D89021.T2205197.A0000015,COPY1VOL=DSNV04
```

Creating a conditional restart control record

To create a new conditional restart control record in the BSDS, you must execute the change log inventory utility and use the CRESTART control statement. For example, to truncate the log, to specify the earliest log RBA, and to bypass back-out, use a statement similar to this:

```
CRESTART CREATE,STARTRBA=28894,ENDRBA=58000,BACKOUT=NO
```

To specify a cold start, make the values of STARTRBA and ENDRBA equal, with a statement similar to this:

```
CRESTART CREATE,STARTRBA=4A000,ENDRBA=4A000
```

In most cases, when doing a cold start, make sure that the STARTRBA and ENDRBA are set to an RBA value greater than the highest RBA used.

An existing conditional restart control record governs any START DB2 operation until one of these events occurs:

- A restart operation completes.
- A CRESTART CANCEL statement is issued.
- A new conditional restart control record is created.

Deleting log data sets with errors

If an active log data set has encountered an I/O error, perform the following steps:

1. If you have been using dual active log data sets, check if the data from the bad active log data set is saved in the other active log. If it is, you can use the other active log.
2. If you cannot use the other active log or the active log is in the STOPPED status, you must fix the problem manually.
   a. Check to see if the data set has been offloaded. For example, check the list of archive log data sets to see if one has the same RBA range as the active log data set. This list can be created by using the DSNJU004 (Print Log Map) utility.
   b. If the data set has not been offloaded, copy the data to a new VSAM data set. If the data set has been offloaded, create a new VSAM data set that is to be used as an active log data set.
   c. Specify DELETE to remove information about the bad data set from the BSDS.

   d. Specify NEWLOG to identify the new data set as the new active log. The DELETE and NEWLOG operations can be performed by the same job step (the DELETE statement precedes the NEWLOG statement in the SYSIN input data set).
3. Delete the bad data set, using VSAM access method services.

Use the Print Log Map utility before and after running the Change Log Inventory utility to ensure correct execution and to document changes.

When using dual active logs, choose a naming convention that distinguishes primary and secondary active log data set. The naming convention should also identify the log data sets within the series of primary or secondary active log data sets. For example, the default naming convention established at DB2 installation time is:

```
prefix.LOGCOPYn.DSmm
```

where *n*=1 for all primary dog data sets and *n*=2 for all secondary log data sets, and *mm* is the data set number within each series.

If a naming convention such as the default convention is used, pairs of data sets with equal mm values are usually used together. For example, DSNC120.LOGCOPY1.DS02 and DSNC120.LOGCOPY2.DS02 are used together.

However, after running the Change Log Inventory utility with the DELETE and NEWLOG statements, the primary and secondary series can become unsynchronized, even if the NEWLOG data set name you specify is the same as the old data set name. To avoid this situation, always do maintenance on both data sets of a pair in the same Change Log Inventory execution:
- Delete both data sets together.
- Define both data sets together with NEWLOG statements.

The data set themselves do not require deletion and redefinition.

To ensure consistent results, execute the Change Log Inventory utility on the same MVS system where the DB2 online subsystem is executing.

If misused, the Change Log Inventory utility can compromise the viability and integrity of the DB2 subsystem. Only highly-skilled people, such as the DB2 System Administrator, should use this utility, and then only after careful consideration.

Before initiating a conditional restart or cold restart, you should consider making backup copies of all DASD volumes containing any DB2 data sets. This will enable a possible fallback. The backup data sets must be generated when DB2 is not active.

Altering references to NEWLOG and DELETE data sets

The NEWLOG and DELETE statements add and delete references to data sets in the BSDS. The log data sets are not changed in any way. If DELETE and NEWLOG are used for a reference in the BSDS to an active log data set, the referenced log data set itself does not require alteration.

Specifying the NEWCAT statement

NEWCAT defines the high-level qualifier used for:
- Catalog table spaces and index spaces
- Directory table spaces and index spaces

At startup, the DB2 system checks that the name recorded with NEWCAT in the BSDS is the high-level qualifier of the DB2 system table spaces that are defined in the load module for subsystem parameters.

NEWCAT is normally used only at installation time. See “Renaming DB2 system data sets” for an additional function of NEWCAT.

When you change the high-level qualifier using the NEWCAT statement, you might specify:

```
//S2 EXEC PGM=DSNJU003
//SYSUT1   DD DSN=DSNC120.BSDS01,DISP=OLD
//SYSUT2   DD DSN=DSNC120.BSDS02,DISP=OLD
//SYSPRINT DD SYSOUT=*
  NEWCAT VSAMCAT=DBP1
```

After running the Change Log Inventory utility with the NEWCAT statement, output similar to the output in Figure 33 is generated.

```
 NEWCAT VSAMCAT=DBP1
SNJ210I OLD VASAM CATALOG NAME=DSNC120 NAME=DBP1
DSNJ225I NEWCAT OPERATION COMPLETED SUCCESSFULLY
DSNJ200I DSNJU003 CHANGE LOG INVENTORY UTILITY
         PROCESSING COMPLETED SUCCESSFULLY
```

Figure 33. Output produced when changing high-level qualifier

Renaming DB2 system data sets

Occasionally, you may want to rename the DB2 system table spaces. In that case you should perform the following steps:

1. Stop DB2 in a consistent state.
2. Create a full system backup so you can recover from operational errors.
3. Execute the Change Log Inventory utility with NEWCAT.
4. Rename the BSDS and all DB2 directory and catalog table spaces and index spaces with IDCAMS.
5. Reassemble DSNZPARM to redefine the high-level qualifier for the system table spaces.
6. Update the BSDS name in the DB2 startup procedure.
7. Start DB2.
8. Drop and recreate the work file database.
9. Optionally use the ALTER command for table spaces in DSNDB04 and user databases.

Renaming DB2 active log data sets

When you rename system data sets, you may also want to rename the log data sets. In that case:

1. Stop DB2 in a consistent state.
2. Create a full system backup so you can recover from operational errors.
3. Delete the reusable active log data sets with IDCAMS, but keep the current active log.
4. Define a new set of active log data sets with IDCAMS.
5. Execute the Change Log Inventory utility to remove names of deleted active log data sets and to define the new active log data set names in the BSDS.

6. Start and use DB2 normally.

When the current active log is archived and becomes reusable, it can be deleted.

Renaming DB2 archive log data sets

You do not need to rename archive log data sets, because:

- Old archive logs are replaced as a part of the normal maintenance cycle.
- The RECOVER utility works with archive logs containing different high-level qualifiers.

To modify the high-level qualifier for archive log data sets, you need to reassemble DSNZPARM.

Sample control statements

Example 1: Adding a new archive log data set

```
NEWLOG DSNAME=DSNREPAL.A0001187,COPY1VOL=DSNV04,UNIT=SYSDA,
STARTRBA=3A190000,ENDRBA=3A1F0000,CATALOG=NO
```

Example 2: Deleting a data set

```
DELETE DSNAME=DSNREPAL.A0001187,COPY1VOL=DSNV04
```

Example 3: Creating a new conditional restart record. The following statement creates a new conditional restart control record, specifying no backward-log recovery and log truncation (a new relative byte address for the end of the log).

```
CRESTART CREATE,BACKOUT=NO,ENDRBA=000000010000
```

Example 4: Adding a communication record to the BSDS

```
DDF LOCATION=USIBMSTODB22,LUNAME=STL#M08,PASSWORD=$STL@290
```

Chapter 33. DSNJU004 (Print Log Map)

The Print Log Map (DSNJU004) utility lists the following information:

- Log data set name, log RBA association, and log LRSN for both copy1 and copy2 of all active and archive log data sets
- Active log data sets that are available for new log data
- Status of all conditional restart control records in the bootstrap data set
- Contents of the queue of checkpoint records in the bootstrap data set
- The communication record of the BSDS, if one exists
- Contents of the quiesce history record
- System and utility timestamps
- Contents of the checkpoint queue

In a data sharing environment, the DSNJU004 utility can list information from any or all BSDSs of a data sharing group.

Additional information regarding the DSNJU004 utility appears in Part 4 (Volume 1) of *DB2 Administration Guide*.

Syntax and options of the control statement

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

Using the SYSIN data set allows you to list information from any or all BSDSs of a data sharing group.

DSNJU004 (print log map) syntax diagram

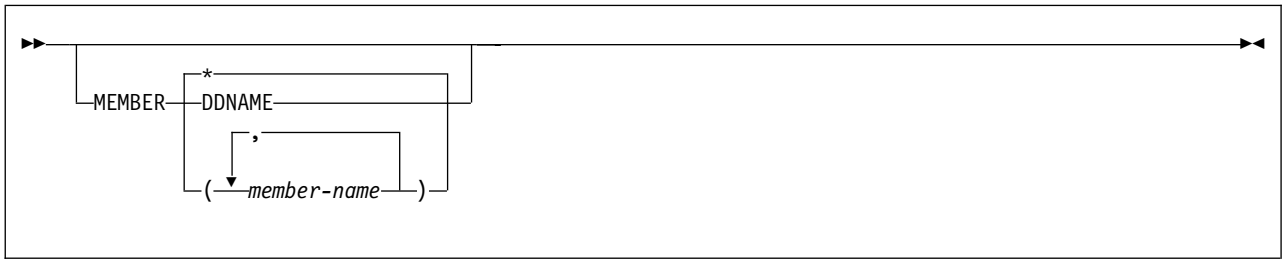

Option descriptions

The following keywords can be used in an optional control statement on the SYSIN data set:

MEMBER

This option specifies which member’s BSDS information to print.

***** Prints the information from the BSDS of each member in the data sharing group.

DDNAME Prints information from only those BSDSs pointed to by the M*xx*BSDS DD statements.

(*member-name***)**

Prints information for only the group members named.

© Copyright IBM Corp. 1983, 2001

Before running DSNJU004

This section contains information you need to be aware of prior to running DSNJU004.

Environment

The DSNJU004 program runs as a batch job.

This utility can be executed both when DB2 is running and when it is not running. However, to ensure consistent results from the utility job, the utility and the DB2 online subsystem must both be executing under the control of the same MVS system.

Authorization required

The user ID of the DSNJU004 job must have requisite RACF authorization.

Control statement

See "DSNJU004 (print log map) syntax diagram" on page 547 for DSNJU004 syntax and option descriptions. See "Sample control statement" on page 549 for an example of a control statement.

Required and optional data sets

DSNJU004 recognizes DD statements with the following ddnames:

JOBCAT

STEPCAT Specifies the catalog in which the bootstrap data set (BSDS) is cataloged. This statement is optional. Typically, the high-level qualifier of the BSDS name points to the integrated catalog facility catalog that contains an entry for the BSDS.

SYSUT1 Specifies and allocates the bootstrap data set. This statement is required. It allocates the BSDS. If the BSDS must be shared with a concurrently executing DB2 online subsystem, use DISP=SHR on the DD statement.

SYSPRINT Specifies a data set or print spool class for print output. This statement is required. The logical record length (LRECL) is 125.

SYSIN (optional)
Contains the control statement. If you do not specify the SYSIN DD statement, BSDS information is printed only from the BSDS data set identified by the SYSUT1 DD statement.

GROUP Names a single BSDS. DB2 can use this BSDS to find the names of all BSDSs in the group. Be sure the BSDS name you specify is not the BSDS of a member that has been quiesced since before new members joined the group. This statement is required if the control statement specifies either of these options:

MEMBER *

MEMBER(*member-name)***)**

M*nn***BSDS** Names the BSDS data set of a group member whose information is to be listed. You must specify one such DD statement for each member. The statements are required if the control statement

specifies **MEMBER DDNAME**. *nn* represents a two-digit number. You must use consecutive two-digit numbers from 01 to the total number of members required. If a break occurs in the sequence of numbers, any number after the break is ignored.

Running the DSNJU004 utility

Use the following EXEC statement to execute this utility:

```
//   EXEC   PGM=DSNJU004
```

Recommendations

- For dual BSDSs, execute the Print Log Map utility twice, once for each BSDS, to compare their contents.
- To ensure consistent results for this utility, execute the utility job on the same MVS system where the DB2 online subsystem is executing.
- Execute the Print Log Map utility regularly, possibly daily, to keep a record of recovery log data set usage.
- Use the Print Log Map utility to document changes made by the Change Log Inventory utility.

Sample control statement

The following statement prints information from the BSDS of each member in the data sharing group:

```
//PLM      EXEC PGM=DSNJU004
//GROUP    DD DSN=DBD1.BSDS01,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
           MEMBER *
```

DSNJU004 (Print Log Map) output

Figure 34 on page 552 shows an example of the Print Log Map utility output, with the following information:

- The data set name (DSN) of the BSDS.
- The system date and time (SYSTEM TIMESTAMP), which is set at the time the subsystem stops.
- The date and time the BSDS was last changed by the Change Log Inventory utility (listed as the UTILITY TIMESTAMP).
- The integrated catalog facility catalog name associated with the BSDS.
- The highest RBA written. The value is updated each time the log buffers are physically written to DASD.
- The highest RBA offloaded.
- Log RBA ranges (STARTRBA and ENDRBA) and data set information for active and archive log data sets. The last active log data set shown is the current active log.
- Information about each active log data set. This information includes the starting and ending RBAs within the data set, the date and time the data set was created, and the data set's name (DSN), and status.
- Information about each archive log data set. This information includes the starting and ending RBAs within the data set, the date and time the data set was created, and the data set's name (DSN), unit and volume of storage, and status.

- Conditional restart control records. For a description of these records and the format of this part of the output from the Print Log Map utility, see “Reading conditional restart control records” on page 554.
- The contents of the checkpoint description queue. For a description of this output, see Figure 35 on page 554.
- The distributed data facility (DDF) communication record. This record contains the DB2-defined location name, the VTAM-defined LUNAME. DB2 uses this information to establish the distributed database environment.

Timestamps in the BSDS

The output of the Print Log Map utility reveals that many timestamps are recorded in the BSDS. Those timestamps record the date and time of various system events.

Timestamps in the output column LTIME are in local time. All other timestamps are in Greenwich Mean Time (GMT).

Figure 34 on page 552 shows an example of the Print Log Map utility output. The following timestamps are included in the header section of the report:

System timestamp
Reflects the date and time the BSDS was last updated. The BSDS can be updated by several events:

- DB2 startup.
- During log write activities, whenever the write threshold is reached.

  Depending on the number of output buffers you have specified and the system activity rate, the BSDS can be updated several times a second, or might not be updated for several seconds, minutes, or even hours.
- When, due to an error, DB2 drops into single-BSDS mode from its normal dual BSDS mode. This can occur when a request to GET, INSERT, POINT to, UPDATE, or DELETE a BSDS record is unsuccessful. When this error occurs, DB2 updates the timestamp in the remaining BSDS to purposely force a timestamp mismatch with the disabled BSDS.

Utility timestamp
The date and time the contents of the BSDS were altered by the Change Log Inventory utility (DSNJU003).

The following timestamps are included in the active and archive log data sets portion of the report:

Active log date
The date the active log data set was originally allocated on the DB2 subsystem.

Active log time
The time the active log data set was originally allocated on the DB2 subsystem.

Archive log date
The date of creation (not allocation) of the archive log data set.

Archive log time
The time of creation (not allocation) of the archive log data set.

The following timestamps are included in the conditional restart control record portion of the report shown in Figure 36 on page 554:

Conditional restart control record
The current time and date. This data is reported as information only and is not kept in the BSDS.

CRCR created The time and date of creation of the CRCR via the CRESTART option in the Change Log Inventory utility.

Begin restart The time and date the conditional restart was attempted.

End restart The time and date the conditional restart ended.

STARTRBA (timestamp) The time the control interval was written.

ENDRBA (timestamp) The time the last control interval was written.

Time of checkpoint The time and date associated with the checkpoint record that was used during the conditional restart process.

The following timestamps are included in the checkpoint queue and the DDF communication record sections of the report shown in Figure 35 on page 554:

Checkpoint queue The current time and date. This data is reported as information only and is not kept in the BSDS.

Time of checkpoint The time and date the checkpoint was taken.

DDF communication record (heading)
The current time and date. This data is reported as information only, and is not kept in the BSDS.

Active log data set status

The BSDS records the status of an active log data set as one of the status values listed in Table 106. For an example of how the status appears in Print Log Map utility output, see Figure 34 on page 552.

Table 106. Statuses of active log data sets

Status	Meaning
NEW	The data set has been defined but never used by DB2, or the log was truncated at a point prior to the data set. In either case, the data set starting and ending RBA values are reset to zero.
REUSABLE	Either the data set is new and has no records, or the data set has been offloaded. In the Print Log Map output, the start RBA value for the last REUSABLE data set is equal to the start RBA value of the last archive log data set.
NOT REUSABLE	The data set contains records that have not been offloaded.
STOPPED	The offload processor encountered an error while reading a record, and that record could not be obtained from the other copy of the active log. Alternatively, an error occurred while truncating the data set following a write I/O error. See Part 4 (Volume 1) of *DB2 Administration Guide*.

Table 106. Statuses of active log data sets (continued)

Status	Meaning
TRUNCATED	One of these conditions exists: • An I/O error occurred, and DB2 has stopped writing to this data set. The active log data set is offloaded, beginning with the starting RBA and continuing up to the last valid record segment in the truncated active log data set. (The RBA of the last valid record segment is less than the ending RBA of the active log data set.) Logging is switched to the next available active log data set and continues uninterrupted. • The log was truncated by a conditional restart at a point within the data set RBA range. • The DB2 ARCHIVE LOG command was issued while this data set was the current active log data set.

```
*************************************************************************************************************
*                                                                                                           *
*              LOG MAP OF THE BSDS DATA SET BELONGING TO MEMBER 'V71A    ' OF GROUP 'DSNCAT  '.             *
*                                                                                                           *
*************************************************************************************************************
RELEASE LEVEL OF BSDS -  ACTIVE=2.3 AND ABOVE   ARCHIVE=2.3 AND ABOVE   DDNAME=SYSUT1
   LOG MAP OF BSDS DATA SET COPY 1, DSN=DSNC710.BSDS01
   LTIME INDICATES LOCAL TIME, ALL OTHER TIMES ARE GMT.
         DATA SHARING MODE IS OFF
         SYSTEM TIMESTAMP   - DATE=2001.057  LTIME= 7:45:29.32
         UTILITY TIMESTAMP  - DATE=2001.057  LTIME= 0:48:16.44
         VSAM CATALOG NAME=DSNC710
         HIGHEST RBA WRITTEN       000001F098DA  2001.057  15:47:16.2
         HIGHEST RBA OFFLOADED     000000000000
         RBA WHEN CONVERTED TO V4  000000000000
       THIS BSDS HAS MEMBER RECORDS FOR THE FOLLOWING MEMBERS:
         HOST MEMBER NAME:            V71A
           MEMBER ID:                 1
           GROUP NAME:                DSNCAT
           BSDS COPY 1 DATA SET NAME: DSNC710.BSDS01
           BSDS COPY 2 DATA SET NAME: DSNC710.BSDS02

ACTIVE LOG COPY 1 DATA SETS
    START RBA/TIME         END RBA/TIME        DATE    LTIME DATA SET INFORMATION
 -------------------  -------------------  --------  ----- --------------------
   000001F02000           000001F05FFF       1998.321  12:48 DSN=DSNC710.LOGCOPY1.DS02
 2001.057  15:46:32.2  2001.057  15:47:03.9  PASSWORD=(NULL) STATUS=TRUNCATED, REUSABLE
   000001F06000           000001F09FFF       1998.321  12:48 DSN=DSNC710.LOGCOPY1.DS03
 2001.057  15:47:03.9  2001.057  15:47:16.2  PASSWORD=(NULL) STATUS=TRUNCATED, REUSABLE
   000001F0A000           00000228DFFF       1998.321  12:48 DSN=DSNC710.LOGCOPY1.DS01
 2001.057  15:47:16.2  ........  ..........  PASSWORD=(NULL) STATUS=REUSABLE
ARCHIVE LOG COPY 1 DATA SETS
NO ARCHIVE DATA SETS DEFINED FOR THIS COPY
ACTIVE LOG COPY 2 DATA SETS
    START RBA/TIME         END RBA/TIME        DATE    LTIME DATA SET INFORMATION
 -------------------  -------------------  --------  ----- --------------------
   000001F02000           000001F05FFF       1998.321  12:48 DSN=DSNC710.LOGCOPY2.DS02
 2001.057  15:46:32.2  2001.057  15:47:03.9  PASSWORD=(NULL) STATUS=TRUNCATED, REUSABLE
   000001F06000           000001F09FFF       1998.321  12:48 DSN=DSNC710.LOGCOPY2.DS03
 2001.057  15:47:03.9  2001.057  15:47:16.2  PASSWORD=(NULL) STATUS=TRUNCATED, REUSABLE
   000001F0A000           00000228DFFF       1998.321  12:48 DSN=DSNC710.LOGCOPY2.DS01
 2001.057  15:47:16.2  ........  ..........  PASSWORD=(NULL) STATUS=REUSABLE
ARCHIVE LOG COPY 2 DATA SETS
NO ARCHIVE DATA SETS DEFINED FOR THIS COPY
DSNJ401I  DSNRJPCR RESTART CONTROL RECORD NOT FOUND
```

Figure 34. Sample print log map utility output (Part 1 of 2)

```
#                          CHECKPOINT QUEUE
#                       15:47:44 FEBRUARY 26, 2001
#       TIME OF CHECKPOINT       15:47:16 FEBRUARY 26, 2001
#       BEGIN CHECKPOINT RBA             000001F0B97A
#       END CHECKPOINT RBA               000001F0D8DA
#       TIME OF CHECKPOINT       15:47:04 FEBRUARY 26, 2001
#       BEGIN CHECKPOINT RBA             000001F0797A
#       END CHECKPOINT RBA               000001F098DA
#
# ⋮
#
#       TIME OF CHECKPOINT       22:25:43 OCTOBER 17, 2000
#       BEGIN CHECKPOINT RBA             000001E31224
#       END CHECKPOINT RBA               000001E35910
#       SHUTDOWN CHECKPOINT
#
# ⋮
#
#       TIME OF CHECKPOINT       03:03:11 NOVEMBER 18, 1998
#       BEGIN CHECKPOINT RBA             00000102BE24
#       END CHECKPOINT RBA               0000010314E0
#       TIME OF CHECKPOINT       03:02:48 NOVEMBER 18, 1998
#       BEGIN CHECKPOINT RBA             0000010020E0
#       END CHECKPOINT RBA               0000010069EA
#                       ARCHIVE LOG COMMAND HISTORY
#                        15:47:44 FEBRUARY 26, 2001
#    DATE          TIME          RBA         MODE   WAIT   TIME
# ------------  ----------  ------------   -------  ----   -----
# FEB 26, 2001  15:47:16.2  000001F09948   QUIESCE  YES    999
# FEB 26, 2001  15:47:03.9  000001F05948   QUIESCE   NO      5 D
# FEB 26, 2001  15:46:32.1  000001F0196E
# FEB 26, 2001  15:45:28.9  000001EFD822
#                  **** DISTRIBUTED DATA FACILITY ****
#                         COMMUNICATION RECORD
#                       15:47:44 FEBRUARY 26, 2001
# LOCATION=STLEC1 LUNAME=SYEC1DB2 PASSWORD=DB2PW1 GENERICLU=(NULL) PORT=NULL RPORT=NULL
# DSNJ200I  DSNJU004 PRINT LOG UTILITY PROCESSING COMPLETED SUCCESSFULLY
```

Figure 34. Sample print log map utility output (Part 2 of 2)

The ARCHIVE LOG COMMAND HISTORY in the output above was created as follows:

- The first entry in the history was created by issuing the ARCHIVE LOG command:

  ```
  -ARCHIVE LOG MODE(QUIESCE) WAIT(YES) TIME(999)
  ```

- The next entry was created by issuing the ARCHIVE LOG command without a time parameter. The D after the time signifies that the default DSNZPARM TIME value (3 seconds) was used.

  ```
  -ARCHIVE LOG MODE(QUIESCE)
  ```

- The last two entries in the history were created by issuing the ARCHIVE LOG command as follows:

  ```
  -ARCHIVE LOG
  ```

- The values in the TIME column of the ARCHIVE LOG COMMAND HISTORY section of the report represent the time the ARCHIVE LOG command was issued. This time value is saved in the BSDS and is converted to printable format at the time the Print Log Map utility is run. Therefore this value, when printed, can differ from other time values that were recorded concurrently. Some time values are converted to printable format when they are recorded, and then they are saved in the BSDS. These printed values remain the same when the printed report is run.

Reading conditional restart control records

In addition to listing information about log records, the Print Log Map utility lists information about each conditional restart control record and each checkpoint description. A sample description of a checkpoint record in the queue is shown in Figure 35.

```
                  CHECKPOINT QUEUE
               13:02:50 MAY 14, 1999
TIME OF CHECKPOINT       13:35:20 MAY 14, 1999
BEGIN CHECKPOINT RBA            000000047C10
END CHECKPOINT RBA              000000048510
TIME OF CHECKPOINT       13:21:49 MAY 14, 1999
BEGIN CHECKPOINT RBA            000000035010
END CHECKPOINT RBA              000000035780
TIME OF CHECKPOINT       13:01:26 MAY 14, 1999
BEGIN CHECKPOINT RBA            000000029000
END CHECKPOINT RBA              0000000297A0
```

Figure 35. Sample print log map description of checkpoints

A sample description of a conditional restart control record is shown in Figure 36.

```
CRCR IDENTIFIER  0001
     USE COUNT   1
     RECORD STATUS
         CRCR NOT ACTIVE
         SUCCESSFUL RESTART
     PROCESSING STATUS
         FORWARD = YES
         BACKOUT = NO
     STARTRBA                         000000028894
     ENDRBA                           000000058000
     EARLIEST REQUESTED RBA           000000027B00
     FIRST LOG RECORD RBA             0000000288B0
     ORIGINAL CHECKPOINT RBA          00000005A390
     NEW CHECKPOINT RBA (CHKPTRBA)    000000047C10
     END CHECKPOINT RBA               000000048510
     CRCR CREATED              11:21:08 MAY 14, 1999
     BEGIN RESTART             11:26:26 MAY 14, 1999
     END RESTART               11:31:38 MAY 14, 1999
     TIME OF CHECKPOINT        10:34:31 MAY 14, 1999
     RESTART PROGRESS                  STARTED     ENDED
                                       =======     =====
         CURRENT STATUS REBUILD          YES        YES
         FORWARD RECOVERY PHASE          YES        YES
         BACKOUT RECOVERY PHASE          YES        YES
```

Figure 36. Sample print log map description of a CRCR

Chapter 34. DSN1CHKR

The DSN1CHKR utility verifies the integrity of DB2 directory and catalog table spaces. DSN1CHKR scans the specified table space for broken links, broken hash chains, and records that are not part of any link or chain.

Use DSN1CHKR on a regular basis to promptly detect any damage to the catalog and directory.

Syntax and options of the control statement

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

DSN1CHKR syntax diagram

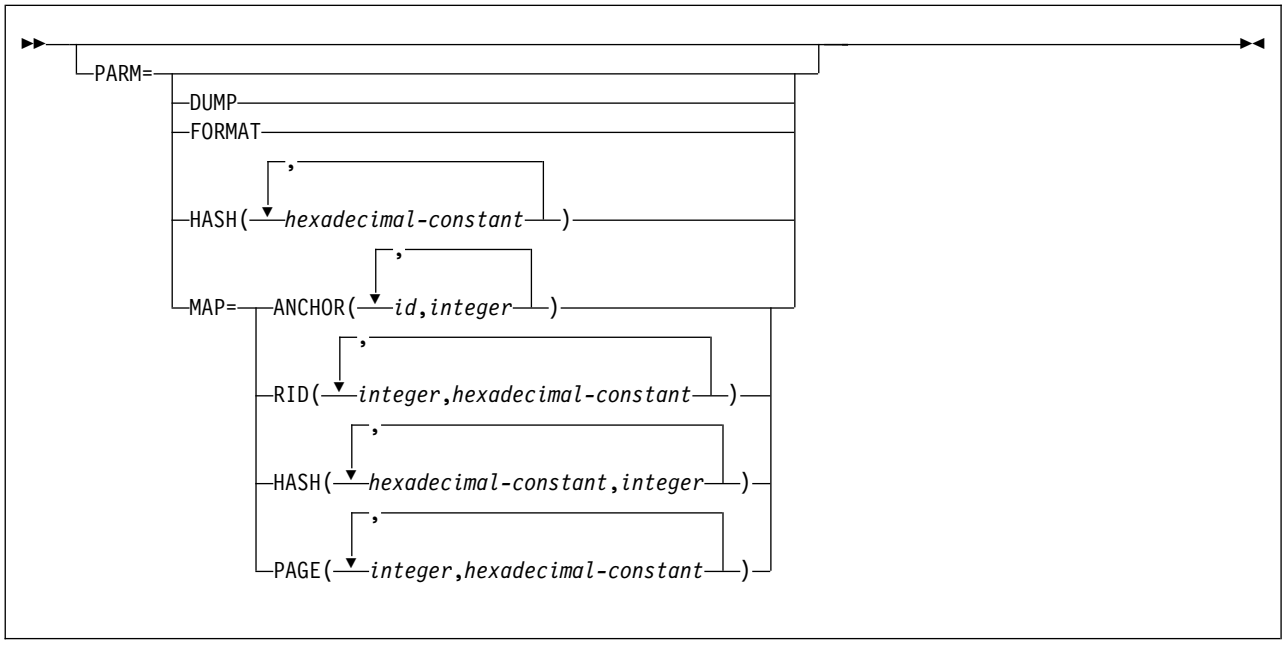

Option descriptions

The following parameters are optional. Parameters should be specified on the EXEC card and can be specified in any order. If you specify more than one parameter, separate them with commas but no blanks. If you do not specify any parameters, DSN1CHKR scans all table space pages for broken links and for records that are not part of any link or chain, and prints the appropriate diagnostic messages.

DUMP Specifies that printed table space pages, if any, are in dump format. If you specify DUMP, you cannot specify the FORMAT parameter.

FORMAT Specifies that printed table space pages, if any, are formatted on output. If you specify FORMAT, you cannot specify the DUMP parameter.

HASH(*hexadecimal-constant***, ...)**
Specifies a hash value for a hexadecimal database identifier (DBID)

© Copyright IBM Corp. 1983, 2001

in table space DBD01. DSN1CHKR returns hash values for each DBID in page and anchor point offset form.

hexadecimal-constant is the hash value for a DBID. The maximum number of DBIDs is 10.

MAP= Identifies a record whose pointer is followed. DSN1CHKR prints each record as it follows the pointer. Use this parameter only after you have determined which chain is broken. You can determine if it is broken by running DSN1CHKR without any parameters, or with FORMAT or DUMP only.

The options for this parameter help DSN1CHKR locate the record whose pointer it follows. Each option must point to the beginning of the 6-byte prefix area of a valid record or to the beginning of the hash anchor. If the value you specify does not point to one of these, DSN1CHKR issues an error message and continues with the next pair of values.

ANCHOR(*id,integer***)**
Specifies the anchor point that DSN1CHKR maps.

id identifies the starting page and anchor point in the form *ppppppaa*, where *pppppp* is the page number and *aa* is the anchor point number.

integer determines which pointer to follow while mapping. 0 specifies the forward pointer; 4 specifies the backward pointer.

The maximum number of pairs is five.

RID(*integer, hexadecimal-constant***, ...)**
Identifies the record or hash anchor from which DSN1CHKR starts mapping.

integer is the page and record, in the form *pppppprr*, where *pppppp* is the page number and *rr* is the record number.

hexadecimal-constant specifies the hexadecimal displacement from the beginning of the record to the pointer in the record from which mapping starts.

The maximum number of pairs is five.

HASH(*hexadecimal-constant, integer***, ...)**
Specifies the value that DSN1CHKR hashes and maps for table space DBD01.

hexadecimal constant is the database identifier in table space DBD01.

integer determines which pointer to follow while mapping. 0 specifies the forward pointer; 4 specifies the backward pointer.

The maximum number of pairs is five.

PAGE(*integer, hexadecimal-constant***, ...)**
integer specifies the page number on which the record or hash anchor is located.

hexadecimal-constant specifies the offset to the pointer from the beginning of the page.

When you use the PAGE option, DSN1CHKR follows the forward pointer while mapping. If a forward pointer does not exist, DSN1CHKR stops mapping after the first record.

The maximum number of pairs is five.

Before running DSN1CHKR

The information in this section is necessary for running DSN1CHKR.

DSN1CHKR is a diagnosis tool; it executes outside the control of DB2. You should have detailed knowledge of DB2 data structures to make proper use of this service aid.

Environment

Run the DSN1CHKR program as a Multiple Virtual Storage (MVS) job.

You must not run DSN1CHKR on a table space while it is active under DB2. While DSN1CHKR runs, do not run other database operations for the database and table space to be checked. Use the STOP DATABASE command for the database and table space to be checked.

Authorization required

This utility does not require authorization. However, if Resource Access Control Facility (RACF) protects any of the data sets, then the authorization must also have the necessary RACF authority.

Control statement

See “Syntax and options of the control statement” on page 555 for DSN1CHKR syntax and option descriptions.

Required data sets: DSN1CHKR uses two data definition (DD) cards. Specify the data set for the utility’s output with the DD card SYSPRINT. Specify the first data set piece of the table space that is to be checked with the DD card SYSUT1.

SYSPRINT Defines the data set that contains output messages from the DSN1CHKR program and all hexadecimal dump output.

SYSUT1 Defines the input data set. This data set can be a DB2 data set or a copy that is created by the DSN1COPY utility. You must specify disposition of this data set as DISP=OLD to ensure that it is not in use by DB2. You must set the data set’s disposition as DISP=SHR only when the STOP DATABASE command has stopped the table space you want to check.

Restrictions

This section contains restrictions to be aware of before running DSN1COMP.

Running DSN1COPY before DSN1CHKR

DSN1CHKR requires a Virtual Storage Access Method (VSAM) data set as input; it cannot check a physical sequential data set.

DSN1CHKR does not use full image copies that are created with the COPY utility. If you create a full image copy with SHRLEVEL REFERENCE, you can copy it into a VSAM data set with DSN1COPY and check it with DSN1CHKR.

DSN1CHKR can not use full image copies that are created with DFSMS concurrent copy. The DFSMS data set does not copy to a VSAM data set because of incompatible formats.

Recommendation: First copy the stopped table space to a temporary data set by using DSN1COPY. Use the DB2 naming convention for the copied data set. Run DSN1CHKR on the copy, which frees the actual table space for restart to DB2.

When you run DSN1COPY, use the CHECK option to examine the table space for page integrity errors. Although DSN1CHKR does check for these errors, running DSN1COPY with CHECK prevents an unnecessary invocation of DSN1CHKR.

Running DSN1CHKR on a valid table space

Run DSN1CHKR only on the following valid table spaces:

- DSNDB01.DBD01
- DSNDB06.SYSDBASE
- DSNDB06.SYSDBAUT
- DSNDB06.SYSGROUP
- DSNDB06.SYSPLAN
- DSNDB06.SYSVIEWS

Sample control statements

Example 1: To run DSN1CHKR on a temporary data set. STEP1 allocates a temporary data set. The fifth qualifier in the data set name can be either I0001 or J0001. Use I0001 in this example. STEP2 stops database DSNDB06 with the STOP DATABASE command. STEP3 copies the target table space into the temporary data set with DSN1COPY. Use the CHECK option to check the table space for page integrity errors. After DSN1COPY with the check option has ensured that no errors exist, STEP4 restarts the table space for access to DB2 again. STEP5 runs DSN1CHKR on the temporary data set.

DSN1CHKR prints the chains, beginning with the pointers on the RID option in the Maintenance Analysis procedure. In this example, the first pointer is on page 2, at an offset of 6 bytes from record 1. The second pointer is on page B, at an offset of 6 bytes from record 1.

The RIDs in STEP5 are for instruction only.

```
//YOUR JOBCARD
//*
//JOBCAT  DD DSNAME=DSNCAT1.USER.CATALOG,DISP=SHR
//STEP1    EXEC PGM=IDCAMS
//*********************************************************************
//*         ALLOCATE A TEMPORARY DATA SET FOR SYSDBASE                *
//*********************************************************************
//SYSPRINT DD SYSOUT=A
//SYSUDUMP DD SYSOUT=A
//SYSIN    DD *
   DELETE                                                    -
        (TESTCAT.DSNDBC.TEMPDB.TMPDBASE.I0001.A001)          -
          CATALOG(DSNCAT)
   DEFINE CLUSTER                                            -
        ( NAME(TESTCAT.DSNDBC.TEMPDB.TMPDBASE.I0001.A001)    -
          NONINDEXED                                         -
          REUSE                                              -
          CONTROLINTERVALSIZE(4096)                          -
          VOLUMES(XTRA02)                                    -
          RECORDS(783 783)                                   -
          RECORDSIZE(4089 4089)                              -
          SHAREOPTIONS(3 3) )                                -
   DATA                                                      -
        ( NAME(TESTCAT.DSNDBD.TEMPDB.TMPDBASE.I0001.A001))   -
          CATALOG(DSNCAT)
/*
//STEP2    EXEC PGM=IKJEFT01,DYNAMNBR=20
//*********************************************************************
//*              STOP DSNDB06.SYSDBASE                                *
//*********************************************************************
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSTSPRT DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSTSIN  DD *
 DSN SYSTEM(V61A)
     -STOP DB(DSNDB06) SPACENAM(SYSDBASE)
 END
/*
//STEP3    EXEC PGM=DSN1COPY,PARM=(CHECK)
//*********************************************************************
//*              CHECK SYSDBASE AND RUN DSN1COPY                      *
//*********************************************************************
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DSN=DSNCAT.DSNDBC.DSNDB06.SYSDBASE.I0001.A001,DISP=SHR
//SYSUT2   DD DSN=TESTCAT.DSNDBC.TEMPDB.TMPDBASE.I0001.A001,DISP=SHR
/*
```

Figure 37. job control language (JCL) for running DSN1CHKR on a temporary data set (Part 1 of 2)

```
//STEP4    EXEC PGM=IKJEFT01,DYNAMNBR=20
//*******************************************************************
//*              START DSNDB06.SYSDBASE                             *
//*******************************************************************
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSTSPRT DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSTSIN  DD *
 DSN SYSTEM(V61A)
     -START DB(DSNDB06) SPACENAM(SYSDBASE)
 END
/*//STEP5    EXEC  PGM=DSN1CHKR,PARM='MAP=RID(00000201,06,00000B01,06)',
//         COND=(4,LT)
//*******************************************************************
//*              CHECK LINKS OF SYSDBASE                            *
//*******************************************************************
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DSN=TESTCAT.DSNDBC.TEMPDB.TMPDBASE.I0001.A001,DISP=SHR
/*
```

Figure 37. job control language (JCL) for running DSN1CHKR on a temporary data set (Part 2 of 2)

Example 2: To Run DSN1CHKR on an actual table space. STEP1 stops database DSNDB06 with the STOP DATABASE command. STEP2 runs DSN1CHKR on the target table space; its output is identical to the output in Example 1. STEP3 restarts the database with the START DATABASE command.

```
//YOUR JOBCARD
//*
//STEP1    EXEC PGM=IKJEFT01,DYNAMNBR=20
//*******************************************************************
//*                        EXAMPLE 2                                *
//*                                                                 *
//*                STOP DSNDB06.SYSDBASE                            *
//*******************************************************************
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSTSPRT DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSTSIN  DD *
 DSN SYSTEM(V61A)
     -STOP DB(DSNDB06) SPACENAM(SYSDBASE)
 END
/*
```

Figure 38. Sample JCL for running DSN1CHKR on a stopped table space. (Part 1 of 2)

```
//STEP2    EXEC  PGM=DSN1CHKR,PARM='MAP=RID(00000201,06,00000B01,06)',
//         COND=(4,LT)
//*******************************************************************
//*              CHECK LINKS OF SYSDBASE                            *
//*******************************************************************
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DSN=DSNCAT.DSNDBD.DSNDB06.SYSDBASE.I0001.A001,DISP=SHR
/*
//STEP3    EXEC PGM=IKJEFT01,DYNAMNBR=20
//*******************************************************************
//*              RESTART DSNDB06.SYSDBASE                           *
//*******************************************************************
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSTSPRT DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSTSIN  DD *
 DSN SYSTEM(V61A)
     -START DB(DSNDB06) SPACENAM(SYSDBASE)
 END
/*
```

Figure 38. Sample JCL for running DSN1CHKR on a stopped table space. (Part 2 of 2)

DSN1CHKR output

Interpreting output

One intended use of this utility is to aid in determining and correcting system problems. When diagnosing DB2, you might need to refer to licensed documentation to interpret output from this utility.

Chapter 35. DSN1COMP

DSN1COMP estimates space savings to be achieved by DB2 data compression in table spaces. For more information regarding ESA data compression, see Part 5 (Volume 2) of *DB2 Administration Guide*.

This utility can be run on the following types of data sets containing uncompressed data:

- DB2 full image copy data sets
- VSAM data sets that contain DB2 table spaces
- Sequential data sets that contain DB2 table spaces (for example, DSN1COPY output)

DSN1COMP does not estimate savings for data sets that contain LOB table spaces or index spaces.

Syntax and options of the control statement

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

DSN1COMP syntax diagram

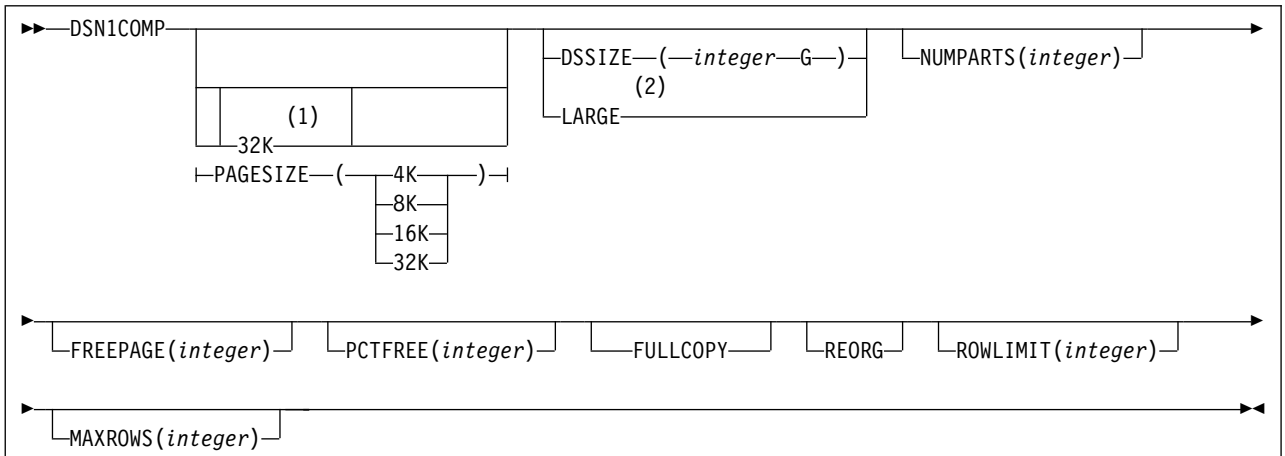

Notes:

1 The preferred option is PAGESIZE(32K).

2 The preferred method of specifying LARGE is DSSIZE(4G).

Important: If you specify more than one parameter:

- Separate them by commas (no blanks).
- Specify them in any order.

Option descriptions

Specify one or more of the parameters listed below on the EXEC card to run DSN1COMP.

© Copyright IBM Corp. 1983, 2001

32K Specifies that the input data set, SYSUT1, has a 32-KB page size. If the SYSUT1 data set has a 32-KB page size, and you do not specify this option, DSN1COMP produces unpredictable results, because the default page size is 4 KB.

The preferred option is **PAGESIZE(32K)**.

PAGESIZE Specifies the page size of the input data set that is defined by SYSUT1. Available page size values are 4K, 8K, 16K, or 32K. If you specify an incorrect page size, DSN1COMP may produce unpredictable results.

If you omit PAGESIZE, DSN1COMP tries to determine the page size from the input data set. DB2 issues an error message if DSN1COMP cannot determine the input page size. This might happen if the header page is not in the input data set, or the page size field in the header page contains an invalid page size.

DSSIZE(*integer* **G)**
Specifies the data set size, in gigabytes, for the input data set. If you omit DSSIZE, DB2 assumes that the input data set size is 2 GB.

integer must match the DSSIZE value specified when the table space was defined.

If you omit DSSIZE and the data set is not one of the default sizes, the results from DSN1COMP are unpredictable.

LARGE Specifies that the input data set is table space that was defined with the LARGE option. If you specify LARGE, then DB2 assumes that the data set has a 4-GB boundary.

The preferred method of specifying a table space defined with LARGE is **DSSIZE(4G)**.

If you omit the LARGE or DSSIZE(4G) option when it is needed, or if you specify LARGE for a table space that was not defined with the LARGE option, the results from DSN1COMP are unpredictable.

NUMPARTS(*integer***)**
Specifies the number of partitions associated with the input data set. Valid specifications range from 1 to 254. If you omit NUMPARTS or specify it as 0, DSN1COMP assumes that your input file is not partitioned. If you specify a number greater than 64, DSN1COMP assumes that the data set is for a partitioned table space that was defined with the LARGE option, even if the LARGE keyword is not specified.

DSN1COMP cannot always validate the NUMPARTS parameter. If you specify it incorrectly, DSN1COMP may produce unpredictable results.

DSN1COMP terminates and issues message DSN1946I when it encounters an image copy containing multiple partitions; a compression report is issued for the first partition.

FREEPAGE(*integer***)**
Specifies how often to leave a page of free space when calculating the percentage of pages saved. You must specify an integer in the range 0 to 255. If you specify 0, no pages are included as free

space when reporting the percentage of pages saved. Otherwise, one free page is included after every *n* pages, where *n* is the specified integer.

The **default** is **0**.

Specify the same value that you specify for the FREEPAGE option of the SQL statement CREATE TABLESPACE or ALTER TABLESPACE.

PCTFREE(*integer***)**
Indicates what percentage of each page to leave as free space when calculating the percentage of pages saved. You must specify an integer in the range 0 to 99. When calculating the savings, DSN1COMP allows for at least *n* percent of free space for each page, where *n* is the specified integer.

The **default** is **5**.

Specify the same value that you specify for the PCTFREE option of the SQL statement CREATE TABLESPACE or ALTER TABLESPACE.

FULLCOPY Specifies that a DB2 full image copy (not a DFSMS concurrent copy) of your data is used as input. Omitting this parameter when the input is a full image copy can cause error messages or unpredictable results. If this data is partitioned, also specify the NUMPARTS parameter to identify the number of partitions.

REORG Provides an estimate of compression savings comparable to the savings that the REORG utility would achieve. If this keyword is not specified, the results are similar to the compression savings that the LOAD utility would achieve.

ROWLIMIT(*integer***)**
Specifies the maximum number of rows to evaluate in order to provide the compression estimate. This option prevents DSN1COMP from examining every row in the input data set. Valid specifications range from 1 to 99,000,000.

Use this option to limit the elapsed time and processor time that DSN1COMP requires. An analysis of the first 5-10 MB of a table space provides a fairly representative sample of the table space for estimating compression savings. Therefore, specify a ROWLIMIT value that restricts DSN1COMP to the first 5-10 MB of the table space. For example, if the row length of the table space is 200 bytes, specifying ROWLIMIT(50000) causes DSN1COMP to analyze approximately 10 MB of the table space.

MAXROWS(*integer***)**
Specifies the maximum number of rows that DSN1COMP is to consider when calculating the percentage of pages saved. You must specify an integer in the range 1 to 255. The **default** is **255**.

Specify the same value that you specify for the MAXROWS option of the SQL statement CREATE TABLESPACE or ALTER TABLESPACE.

Before running DSN1COMP

If you run DSN1COMP on a segmented table space, you must first query DB2 catalog table, SYSTABLEPART, to determine the current instance qualifier. Use the current instance qualifier in the job control language (JCL).

```
SELECT DBNAME, TSNAME, PARTITION, IPREFIX FROM SYSIBM.SYSTABLEPART
WHERE DBNAME = 'DBDV0702' AND TSNAME  = 'TPDV0702';
```

The above query would produce the following example information.

DBNAME	TSNAME	PARTITION	IPREFIX
DBDV0702	TPDV0702	0	J

The example output of the above query would provide the user with the information needed to code the data set name in the DSN1COMP JCL.

```
//STEP1 EXEC PGM=DSN1COMP
//SYSUT1 DD DSN=vcatname.DSNDBC.DBDV0702.J0001.A001,DISP=SHR
//SYSPRINT DD AYAOUT=*
//SYSUDUMP DD AYAOUT=*
```

Environment

Run DSN1COMP as a Multiple Virtual Storage (MVS) job.

You can run DSN1COMP even when the DB2 subsystem is not operational. Before you use DSN1COMP when the DB2 subsystem is operational, issue the DB2 STOP DATABASE command. The STOP DATABASE command will ensure that DB2 has not allocated the DB2 data sets.

You must not run DSN1COMP on table spaces in DSNDB01, DSNDB06, or DSNDB07.

Authorization required

DSN1COMP does not require authorization. However, if any of the data sets is Resource Access Control Facility (RACF)-protected, the authorization ID of the job must have RACF authority.

Control statement

See “Syntax and options of the control statement” on page 563 for DSN1COMP syntax and option descriptions.

Required data sets: DSN1COMP uses the data definition (DD) cards described below:

SYSPRINT Defines the data set that contains output messages from DSN1COMP and all hexadecimal dump output.

SYSUT1 Defines the input data set, which can be a sequential data set or a Virtual Storage Access Method (VSAM) data set.

You must specify the disposition for this data set as OLD (DISP=OLD) to ensure that it is not in use by DB2. You must specify the disposition for this data set as SHR (DISP=SHR) only in circumstances where the DB2 STOP DATABASE command does not work.

The requested operation takes place only for the data set specified. The following are instances when you must specify the correct data set.

- The input data set belongs to a linear table space
- The index space is larger than 2 gigabytes
- The table space or index space is a partitioned space

Recommendation

This section contains a recommendation for running DSN1COMP.

Determination of page size and DSSIZE: Before using DSN1COMP, be sure you know the page size and data set size (DSSIZE) for the table space. Use the following query on the DB2 catalog to get the information you need:

```
SELECT  T.CREATOR,
        T.NAME,
        S.PGSIZE,
        CASE S.DSSIZE
          WHEN 0 THEN
            CASE S.TYPE
              WHEN ' ' THEN 2097152
              WHEN 'I' THEN 2097152
              WHEN 'L' THEN 4194304
              WHEN 'K' THEN 4194304
              ELSE NULL
            END
          ELSE S.DSSIZE
        END
  FROM  SYSIBM.SYSTABLES  T,
        SYSIBM.SYSTABLESPACE S
 WHERE  T.DBNAME=S.DBNAME
   AND  T.TSNAME=S.NAME;
```

Using DSN1COMP

This section describes the following tasks associated with running the DSN1COMP utility:

"Estimating compression savings achieved by REORG"
"Including free space in compression calculations" on page 568
"Running DSN1COMP on a table space with identical data" on page 568

Estimating compression savings achieved by REORG

If you run DSN1COMP with the REORG option on small data sets, the estimates produced might vary greatly from the estimates produced without the default REORG option. Or, if you run DSN1COMP and specify a small number (*n*) for ROWLIMIT, the estimates might vary greatly from the estimates produced without REORG.

Without the REORG option, DSN1COMP uses the first *n* rows to fill the compression dictionary. DSN1COMP processes the remaining rows to provide the compression estimate. If the number of rows that are used to build the dictionary is a significant percentage of the data set rows, little savings will result. With the REORG option, DSN1COMP processes all the rows, including those that are used to build the dictionary, with greater compression.

Including free space in compression calculations

DSN1COMP utility's compression estimates take into account the PCTFREE and FREEPAGE options. If you use different PCTFREE or FREEPAGE values than those created with the input table space, you get a different value for **noncmppages**. DSN1COMP reports this in message DSN1940I.

Running DSN1COMP on a table space with identical data

If you run DSN1COMP on a table space in which the data is the same for all rows, you will receive message DSN1941I. In this case, DSN1COMP does not compute any statistics.

Sample control statements

Example 1: To run DSN1COMP

The fifth qualifier in the data set name can be either I0001 or J0001. In this example we are using I0001.

```
//jobname JOB acct information
//COMPEST  EXEC PGM=DSN1COMP,PARM='FULLCOPY'
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSABEND DD SYSOUT=A
//SYSUT1   DD DSN=DSNCAT.DSNDBC.DB254A.TS254A.I0001.A001,DISP=SHR
```

Example 2: To run DSN1COMP with the PCTFREE and FREEPAGE options

The fifth qualifier in the data set name can be either I0001 or J0001. In this example we are using I0001.

```
//DSN1COMP JOB  MSGLEVEL=(1,1),CLASS=A,MSGCLASS=A,REGION=3000K,
//         USER=SYSADM,PASSWORD=SYSADM
/*ROUTE PRINT STLXXXX.USERID
//STEP1    EXEC PGM=DSN1COMP,PARM='PCTFREE(20),FREEPAGE(5)'
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSDUMP  DD SYSOUT=A
//SYSABEND DD SYSOUT=A
//SYSUT1   DD DSN=DSNC710.DSNDBD.DB254SP4.TS254SP4.I0001.A001,DISP=SHR
/*
//STEP2    EXEC PGM=DSN1COMP,PARM='ROWLIMIT(20000)'
//STEPLIB  DD DSN=prefix.SDSNLOAD,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSDUMP  DD SYSOUT=A
//SYSABEND DD SYSOUT=A
//SYSUT1   DD DSN=DSNC710.DSNDBD.DB254SP4.TS254SP4.I0001.A001,DISP=SHR
/*
//
```

DSN1COMP output

This section contains examples of output that is generated by the DSN1COMP utility.

Message DSN1941

If you receive this message, use a data set with more rows as input, or specify a larger ROWLIMIT.

Sample DSN1COMP report

Figure 39 shows a sample of the output that DSN1COMP generates.

```
DSN1940I DSN1COMP COMPRESSION REPORT
             301  KB WITHOUT COMPRESSION
             224  KB WITH COMPRESSION
              25  PERCENT OF THE BYTES WOULD BE SAVED

           1,975  ROWS SCANNED TO BUILD DICTIONARY
           4,665  ROWS SCANNED TO PROVIDE COMPRESSION ESTIMATE
           4,096  DICTIONARY ENTRIES

              81  BYTES FOR AVERAGE UNCOMPRESSED ROW LENGTH
              52  BYTES FOR AVERAGE COMPRESSED ROW LENGTH

              16  DICTIONARY PAGES REQUIRED
             110  PAGES REQUIRED WITHOUT COMPRESSION
              99  PAGES REQUIRED WITH COMPRESSION
              10  PERCENT OF THE DB2 DATA PAGES WOULD BE SAVED
```

Figure 39. Sample DSN1COMP report

Chapter 36. DSN1COPY

With the DSN1COPY stand-alone utility, you can copy:

- DB2 VSAM data sets to sequential data sets
- DSN1COPY sequential data sets to DB2 VSAM data sets
- DB2 image copy data sets to DB2 VSAM data sets
- DB2 VSAM data sets to other DB2 VSAM data sets
- DSN1COPY sequential data sets to other sequential data sets

Using DSN1COPY, you can also print hexadecimal dumps of DB2 data sets and databases, check the validity of data or index pages (including dictionary pages for compressed data), translate database object identifiers (OBIDs) to enable moving data sets between different systems, and reset to 0 the log RBA that is recorded in each index page or data page.

DSN1COPY is compatible with LOB table spaces, when you specify the LOB keyword, and omit the SEGMENT and INLCOPY keywords.

Syntax and options of the control statement

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

© Copyright IBM Corp. 1983, 2001

DSN1COPY syntax diagram

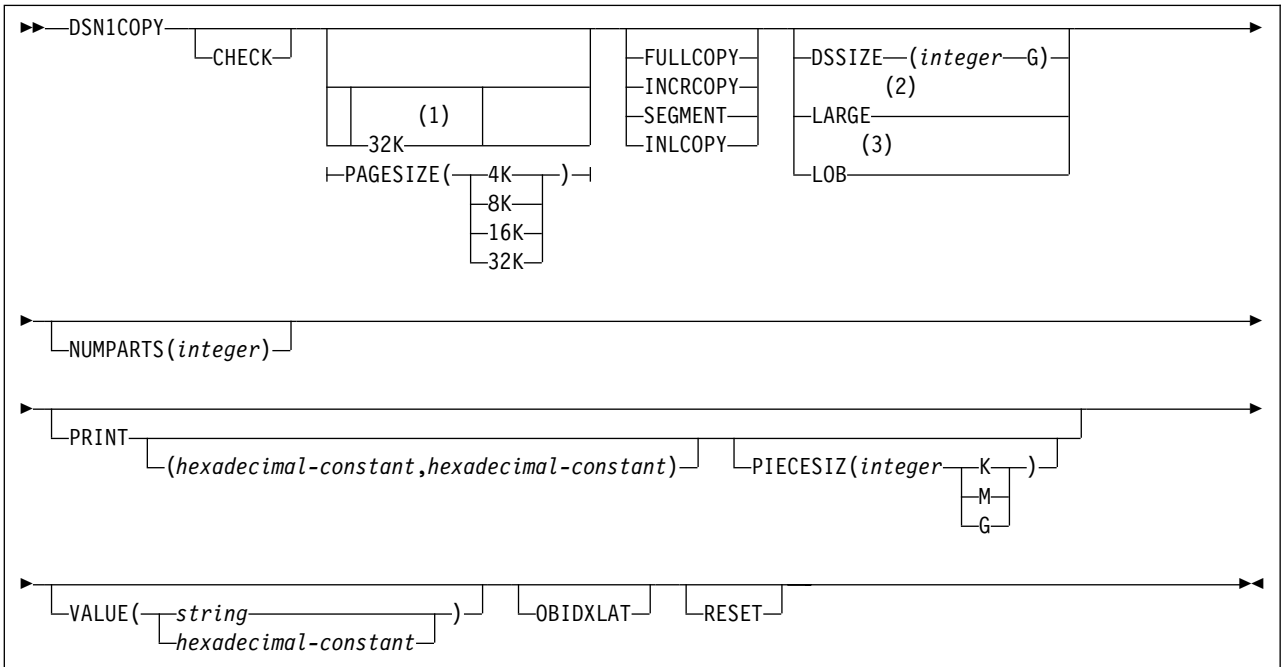

Notes:

1 The preferred option is PAGESIZE(32K).

2 The preferred method of specifying LARGE is DSSIZE(4G).

3 You cannot specify the SEGMENT and INLCOPY keywords with the LOB keyword.

Important: If you specify more than one parameter:
- Separate them by commas (no blanks).
- Specify them in any order.

Option descriptions

Specify one or more of the parameters listed below on the EXEC card to run DSN1COPY.

CHECK Checks each page from the SYSUT1 data set for validity. The validity checking operates on one page at a time and does not include any cross-page checking. If an error is found, a message is issued describing the type of error, and a dump of the page is sent to the SYSPRINT data set. If you do not receive any messages, no errors were found. If more than one error exists in a given page, the check identifies only the first of the errors. However, the entire page is dumped.

32K Specifies that the SYSUT1 data set has a 32-KB page size. If the SYSUT1 data set has a 32-KB page size, and you do not specify this option, DSN1COPY may produce unpredictable results, because the default page size is 4 KB.

The preferred option is **PAGESIZE(32K)**.

PAGESIZE Specifies the page size of the input data set that is defined by SYSUT1. Available page size values are 4K, 8K, 16K, or 32K. If you specify an incorrect page size, DSN1COPY may produce unpredictable results.

If you omit PAGESIZE, DSN1COPY tries to determine the page size from the input data set. DB2 issues an error message if DSN1COPY cannot determine the input page size. This might happen if the header page is not in the input data set, or if the page size field in the header page contains an invalid page size.

FULLCOPY Specifies that a DB2 full image copy (not a DFSMS concurrent copy) of your data is to be used as input. If this data is partitioned, specify NUMPARTS to identify the total number of partitions. If you specify FULLCOPY without NUMPARTS, DSN1COPY assumes that your input file is not partitioned.

Specify FULLCOPY when using a full image copy as input. Omitting the parameter can cause error messages or unpredictable results.

The FULLCOPY parameter requires SYSUT2 (output data set) to be either a DB2 VSAM data set or a DUMMY data set.

INCRCOPY Specifies that an incremental image copy of the data is used as input. DSN1COPY with the INCRCOPY parameter updates existing data sets; do not redefine the existing data sets. INCRCOPY requires that the output data set (SYSUT2) be a DB2 VSAM data set.

Before you apply an incremental image copy to your data set, you must first apply a full image copy to the data set using the FULLCOPY parameter. Make sure that you apply the full image copy in a separate execution step, because you receive an error message if you specify both the FULLCOPY and the INCRCOPY parameters in the same step. Then, apply each incremental image copy in a separate step starting with the oldest incremental image copy.

Specifying neither FULLCOPY nor INCRCOPY implies that the input is not image copy data sets. Therefore, only a single output data set is used.

SEGMENT Specifies that you want to use a segmented table space as input to DSN1COPY. Zeroed pages in the table space are copied, but no error messages are issued. You cannot specify FULLCOPY or INCRCOPY if you specify SEGMENT.

If you are using DSN1COPY with the OBIDXLAT to copy a DB2 data set to another DB2 data set, the source and target table spaces must have the same SEGSIZE attribute.

You cannot specify the SEGMENT option with the LOB parameter.

INLCOPY Specifies that the input data is an inline copy data set.

You cannot specify the INLCOPY option with the LOB parameter.

DSSIZE(*integer* **G)**
Specifies the data set size, in gigabytes, for the input data set. If you omit DSSIZE, DB2 assumes that the input data set size is 2 GB unless the input data set is a LOB, in which case DB2 assumes a 4 GB input data set size.

integer must match the DSSIZE value specified when the table space was defined.

If you omit DSSIZE and the data set is not one of the default sizes, the results from DSN1COPY are unpredictable.

LARGE Specifies that the input data set is a table space that was defined with the LARGE option, or an index on such a table space. If you specify the LARGE keyword, DB2 assumes that the data set has a 4-GB boundary.

The preferred method of specifying a table space that was defined with the LARGE option is **DSSIZE(4G)**.

If you omit the LARGE or DSSIZE(4G) option when it is needed, or if you specify LARGE for a table space that was not defined with the LARGE option, the results from DSN1COPY are unpredictable.

LOB Specifies that SYSUT1 data set is a LOB table space. Empty pages in the table space are copied, but no error messages are issued. You cannot specify the SEGMENT and INLCOPY options with the LOB parameter.

DB2 attempts to determine if the input data set is a LOB data set. If you specify the LOB option but the data set is not a LOB data set, or if you omit the LOB option but the data set is a LOB data set, DB2 issues an error message and terminates.

NUMPARTS(*integer***)**
Specifies the total number of partitions associated with the data set you are using as input or whose page range you are printing.

integer can range from 1 to 254.

DSN1COPY uses this value to calculate the size of its output data sets and to help locate the first page in a range to be printed. If you omit NUMPARTS or specify it as 0, DSN1COPY assumes that your input file is not partitioned. If you specify a number greater than 64, DSN1COPY assumes the data set is for a partitioned table space that was defined with the LARGE option, even if the LARGE keyword is not specified for DSN1COPY.

If you specify the number of partitions incorrectly, DSN1COPY can copy the data to the wrong data sets, return an error message indicating that an unexpected page number was encountered, or fail to allocate the data sets correctly. In the last case, a VSAM PUT error might be detected, resulting in a request parameter list (RPL) error code of 24.

PRINT(*hexadecimal-constant,hexadecimal-constant***)**
Causes the SYSUT1 data set to be printed in hexadecimal format on the SYSPRINT data set. You can specify the PRINT parameter with or without the page range specifications (*hexadecimal-constant,hexadecimal-constant*). If you do not specify a range, all pages of the SYSUT1 are printed. If you want to limit the range of pages printed, indicate the beginning and ending page. If you want to print a single page, supply only that page number. In either case, your range specifications must be from one to eight hexadecimal characters in length.

The following example shows how you code the PRINT parameter if you want to begin printing at page X'2F0' and stop at page X'35C':

```
PRINT(2F0,35C)
```

Because the CHECK and RESET options and the COPY function run independently of the PRINT range, these options apply to the entire input file regardless of whether a range of pages is being printed.

PIECESIZ(*integer***)**
Specifies the maximum piece size (data set size) for non-partitioned indexes. The value you specify must match the value specified when the nonpartitioning index was created or altered.

The defaults for PIECESIZ are 2G (2 GB) for indexes backed by non-large table spaces and 4G (4 GB) for indexes backed by table spaces that were defined with the LARGE option. This option is required if a print range is specified and the piece size is not one of the default values. If PIECESIZ is omitted and the index is backed by a table space that was defined with the LARGE option, the LARGE option is required for DSN1COPY.

The subsequent keyword K, M, or G, indicates the units of the value specified in *integer*.

K Indicates that the *integer* value is to be multiplied by 1 KB to specify the maximum piece size in bytes. *integer* must be either 256 or 512.

M Indicates that the *integer* value is to be multiplied by 1 MB to specify the maximum piece size in bytes. *integer* must be a power of two, between 1 and 512.

G Indicates that the *integer* value is to be multiplied by 1 GB to specify the maximum piece size in bytes. *integer* must be 1, 2, or 4.

Valid values for piece size are:
- 1 MB or 1 GB
- 2 MB or 2 GB
- 4 MB or 4 GB
- 8 MB
- 16 MB
- 32 MB
- 64 MB
- 128 MB
- 256 KB or 256 MB
- 512 KB or 512 MB

VALUE Causes each page of the SYSUT1 input data set to be scanned for the character string you specify in parentheses following the VALUE parameter. Each page that contains that character string is printed in the SYSPRINT data set. You can specify the VALUE parameter in conjunction with any of the other DSN1COPY parameters.

string can consist of 1 to 20 alphanumeric characters.

hexadecimal-constant can consist of 2 to 40 hexadecimal characters. You must specify two single quotation mark characters before and after the hexadecimal character string.

If you want to search your input file for the string '12345', your JCL should look like this:

```
//STEP1 EXEC PGM=DSN1COPY,PARM='VALUE(12345)'
```

If you want to search for the equivalent hexadecimal character string, your JCL should look like this:

```
//STEP1 EXEC PGM=DSN1COPY,PARM='VALUE(''F1F2F3F4F5'')'
```

OBIDXLAT Specifies that OBID translation must be done before the DB2 data set is copied. This parameter requires additional input from the SYSXLAT file by using the DD cards. DSN1COPY can only translate up to 500 record OBIDs at a time. If you specify OBIDXLAT, CHECK processing is performed regardless of whether you specify the CHECK option.

RESET Causes the high formatted page number in the header page and the log RBAs in each index or data page to be reset to 0. If you specify this option, CHECK processing is performed regardless of whether you specify the CHECK option.

You must use RESET when the output file is used to build a DB2 table space to be processed on a DB2 subsystem with a different recovery log than the source subsystem. Failure to specify RESET in such a case can result in an abend during subsequent update activity. The abend reason code of 00C200C1 indicates that the specified RBA value is outside the valid range of the recovery log. A condition code of 0 indicates successful completion.

If you do not specify RESET when copying a table space from one DB2 system to another, a down-level ID check may result in abend reason code 00C2010D when the table space is accessed. For more information about down-level detection, see Part 4 (Volume 1) of *DB2 Administration Guide*.

Before running DSN1COPY

This section contains information to use before running DSN1COPY.

Attention: Do not use DSN1COPY in place of COPY or both backup and recovery. Improper use of DSN1COPY can result in unrecoverable damage and loss of data.

Environment

Execute DSN1COPY as a Multiple Virtual Storage (MVS) job, when the DB2 subsystem is either active or not active.

If you execute DSN1COPY when DB2 is active, follow the procedure below:

1. Start the table space as read-only by using START DATABASE.
2. Run the QUIESCE utility with the WRITE (YES) option to externalize all data pages and index pages.
3. Run DSN1COPY with DISP=SHR on the data definition (DD) card.
4. Start the table space as read-write by using -START DATABASE to return to normal operations.

Authorization required

DSN1COPY does not require authorization. However, any data sets protected by Resource Access Control Facility (RACF) require RACF authorization.

Control statement

See “Syntax and options of the control statement” on page 571 for DSN1COPY syntax and option descriptions.

Required data sets

DSN1COPY uses the data sets described below:

Input data set	Input to DSN1COPY. The DD name is SYSUT1.
Output data set	Output from DSN1COPY. The DD name is SYSUT2. Optional.
Message data set	Data set for output messages. The DD name is SYSPRINT.
OBIDXLAT data set	Data set that defines the OBID translation values. The DD name is SYSXLAT.

DSN1COPY uses several DD cards. They are:

SYSPRINT Defines the data set that contains output messages from the DSN1COPY program and all hexadecimal dump output.

SYSUT1 Defines the input data set. This data set can be a sequential data set created by the DSN1COPY or COPY utilities, or a Virtual Storage Access Method (VSAM) data set. DSN1COPY assumes that the block size is 4096 bytes (the standard for DB2 data sets).

You must specify the data set’s disposition as DISP=OLD to ensure that it is not in use by DB2. You must specify the data set’s disposition as DISP=SHR only when the DB2 STOP DATABASE command does not work.

The requested operation takes place only for the specified data set. If the input data set is a partitioned table space or index, then ensure that you specify the NUMPARTS parameter and the correct data set. For example, to print a page range in the second partition of a four-partition table space, specify NUMPARTS(4) and the data set name of the second data set. This second data set is in the group of VSAM data sets. In VSAM data sets it comprises the table space (DSN=...A002).

| If running the online REORG utility with the FASTSWITCH option,
verify the data set name before running the DSN1COPY utility. The
default for the FASTSWITCH parameter is YES. If you do not
specify FASTSWITCH NO, the fifth level qualifier in the data set
name alternates between I0001 and J0001.You must specify the
| correct fifth level qualifier in the data set name to successfully
| execute the DSN1COPY utility.

SYSUT2 Defines the output data set. This data set can be a sequential data set, a VSAM data set, or a DUMMY data set.

DSN1COPY assumes that the output data sets are empty (that is, the program adds the blocks) except when you specify INCRCOPY. Before you run DSN1COPY, define your VSAM output data sets as REUSE. If you have not defined the data sets, you must redefine all VSAM output data sets you are restoring by using Access Method Services. Be sure that these data sets are empty before you run DSN1COPY .

You might want to specify a DUMMY SYSUT2 DD card if you are dumping or checking pages.

To enable DB2 to obtain necessary information from the integrated catalog facility catalog when using VSAM data sets, do not code unit-serial parameter and volume-serial parameter.

If running the online REORG utility with the FASTSWITCH option, verify the data set name before running the DSN1COPY utility. The default for the FASTSWITCH parameter is YES. If you do not specify FASTSWITCH NO, the fifth level qualifier in the data set name alternates between I0001 and J0001.You must specify the correct fifth level qualifier in the data set name to successfully execute the DSN1COPY utility.

SYSXLAT SYSXLAT defines for translation the DBIDs, OBIDs, presentation space ID (PSID), or ISOBIDs.

If you have dropped a table without a subsequent REORG of the table space, you must reorganize the source table space before running DSN1COPY with the OBIDXLAT option. This removes any previously-dropped records from the table space.

A nonnumeric character must separate each record in the SYSXLAT file and each record must contain a pair of decimal integers. The first integer of each record pertains to the source, and the second integer pertains to the target. The first record in the SYSXLAT file contains the source DBIDs and the target DBIDs; the values can range from -32767 to 65535. The second record contains the source and target PSIDs or ISOBIDs; the values can range from 0 to 32767. All subsequent records in the SYSXLAT data set are for table OBIDs. For an index, the SYSXLAT data set must contain the index fan set OBID in addition to the DBID and ISOBID. Sample data in a SYSXLAT file can look like this:

```
260,280
2,10
3,55
6,56
7,57
```

To obtain the names, DBIDs, PSIDs, ISOBIDs, and OBIDs, run the DSNTEP2 sample application on both the source and target systems. The following SQL statements yield the above information:

Note: The example for indexes yields the above information, along with an additional column of data.

Product-sensitive Programming Interface

For table spaces:

```
SELECT DBID, PSID FROM SYSIBM.SYSTABLESPACE
   WHERE NAME='tablespace_name'
     AND DBNAME='database_name';
SELECT NAME, OBID FROM SYSIBM.SYSTABLES
   WHERE TSNAME='tablespace_name'
     AND CREATOR='creator_name';
```

For index spaces:

```
SELECT DBID, ISOBID, OBID FROM SYSIBM.SYSINDEXES
   WHERE NAME='index_name'
     AND CREATOR='creator_name';
```

End of Product-sensitive Programming Interface

Below are several examples of using DSN1COPY:

- Create a backup copy of a DB2 data set:
  - SYSUT1: DB2-VSAM
  - SYSUT2: Sequential data set
- Restore a backup copy of a DB2 data set:
  - SYSUT1: DSN1COPY sequential data set
  - SYSUT2: DB2-VSAM
- Move a DB2 data set to another DB2 data set:
  - SYSUT1: DB2-VSAM
  - SYSUT2: DB2-VSAM
  - Parameters: OBIDXLAT, RESET
- Perform validity checking on a DB2 data set:
  - SYSUT1: DB2-VSAM
  - SYSUT2: DUMMY
  - Parameter: CHECK
- Perform validity checking on and print a DB2 data set:
  - SYSUT1: DB2-VSAM
  - SYSUT2: DUMMY
  - Parameters: CHECK, PRINT
- Restore a table space from a nonpartitioned image copy data set or page set:
  - SYSUT1: DB2 full image copy
  - SYSUT2: DB2-VSAM
  - Parameter: FULLCOPY
- Restore a table space from a partitioned image copy data or page set:
  - SYSUT1: DB2 full image copy
  - SYSUT2: DB2-VSAM
  - Parameters: FULLCOPY, NUMPARTS(*nn*)
- Perform RBA RESET on a DB2 data set:
  - SYSUT1: DB2-VSAM or DSN1COPY sequential data set
  - SYSUT2: DB2-VSAM
  - Parameter: RESET

Defining the input data set

The SYSUT1 data set can be any of the following types:

- A DB2 table space data set
- A DB2 index space data set
- A full image copy
- An incremental image copy
- A sequential data set previously created by DSN1COPY

You should define SYSUT1 with DISP=OLD to ensure that DSN1COPY uses it exclusively. If SYSUT1 is a table space or index space, use the following procedure before using DSN1COPY:

1. Issue the following command to determine if it is a stopped object :

   ```
   -DISPLAY DATABASE (database_name) SPACENAM(space_name) RESTRICT
   ```

2. If DB2 has not stopped the object, issue the following command to stop the object:

```
-STOP DATABASE (database_name) SPACENAME(space_name)
```

DB2 allows input of only one DSN1COPY data set. DB2 does not permit the input of concatenated data sets. For a table space consisting of multiple data sets, ensure that you specify the correct data set. For example, if you specify the CHECK option to validate pages of a partitioned table space's second partition, code the second data set of the table space for SYSUT1.

Defining the output data set

The SYSUT2 data set can be any of the following types:

- A sequential data set
- A DB2 table space data set
- A DB2 index space data set
- A DUMMY data set

Specify a DUMMY SYSUT2 DD card if you are using DSN1COPY to check or dump a page. The table spaces, and index spaces, must either be empty or defined with VSAM REUSE. STOGROUP-defined table spaces and index spaces have the REUSE attribute, except when you are applying the INCRCOPY option

Naming the output data set

For your output data set to be useful, you must make sure that it has the same name as the data set you are resetting.

- Method 1:
  1. Use DSN1COPY to copy your existing data set to a sequential data set. Specify this target data set as SYSUT1.
  2. If you defined your existing data set without the REUSE parameter, delete and redefine the data set. Specify your existing data set as SYSUT2.
- Method 2:
  1. Use your existing DB2 data set as the SYSUT1 specification, creating a new VSAM data set for SYSUT2.
  2. After completion of the reset operation, delete the data set you specified as SYSUT1, and rename the SYSUT2 data set. Give SYSUT2 the name of the data set that you just deleted.

If you are using full or incremental copies as input, specify the SYSUT2 data sets according to the following guidelines:

- **If SYSUT1 is an image copy of a single partition**, SYSUT2 must list the data set name for that partition of the table space. Specify the NUMPARTS parameter to identify the number of partitions in the whole table space.
- **If SYSUT1 is an image copy of a whole partitioned table space**, SYSUT2 must list the name of the table space's first data set. **Attention:** All data sets in the partitioned table space must use the same fifth level qualifier, I0001 or J0001, before DSN1COPY can run successfully on a partitioned table space. DSN1COPY allocates all of the target data sets. However, you must previously define the target data sets by using IDCAMS. Specify the NUMPARTS parameter to identify the number of partitions in the whole table space.
- **If SYSUT1 is an image copy of a nonpartitioned data set**, SYSUT2 should be the name of the actual output data set. Do not specify the NUMPARTS parameter, because this parameter is only for partitioned table spaces.

- **If SYSUT1 is an image copy of all data sets in a linear table space with multiple data sets** , SYSUT2 should be the name of its first data set. DSN1COPY allocates all target data sets. However, you must previously define the target data sets by using IDCAMS.

Adding additional volumes for SYSUT2

When you create a table space or index space by using STOGROUP, the integrated catalog facility catalog entry has only one volume in the volume list. If the SYSUT2 data that DSN1COPY restores requires more than one volume, use the IDCAMS command, ALTER ADDVOLUMES, to add additional volume IDs to the integrated catalog entry. The extension to new volumes uses the primary size on each new volume. This is normal VSAM extension. If you want the data set to use the secondary size on the candidate volumes, do the following steps:

1. Run DSN1COPY.
2. Run REORG, or make a full image copy and recover the table space.

This resets the data set and causes normal extensions through DB2.

Restrictions

This section contains restrictions for running DSN1COPY.

DSN1COPY does not alter data set structure. For example, DSN1COPY does not copy a partitioned or segmented table space into a simple table space. The output data set is a page-for-page copy of the input data set. If the intended use of DSN1COPY is to move or restore data, ensure that definitions for the source and target table spaces, tables, and indexes are identical. Otherwise, unpredictable results can occur.

DSN1COPY cannot copy DB2 recovery log data sets. The format of a DB2 log page is different from that of a table or index page. If you try to use DSN1COPY to recover log data sets, DSN1COPY will abend.

Recommendations

This section contains recommendations for running the DSN1COPY utility.

Printing with DSN1PRNT instead of DSN1COPY

If you require only a printed hexadecimal dump of a data set, use DSN1PRNT rather than DSN1COPY. For more information, see “Printing with DSN1PRNT instead of DSN1COPY” on page 615.

Determining page size and DSSIZE

Before using DSN1COPY, be sure you know the page size and data set size (DSSIZE) for the page set. Use the following query on the DB2 catalog to get the information you need:

```
SELECT I.CREATOR,
      I.NAME,
      S.PGSIZE,
      CASE S.DSSIZE
        WHEN 0 THEN CASE S.TYPE
                      WHEN ' ' THEN 2097152
                      WHEN 'I' THEN 2097152
                      WHEN 'L' THEN 4194304
                      WHEN 'K' THEN 4194304
                      ELSE NULL
                    END
        ELSE S.DSSIZE
      END
 FROM SYSIBM.SYSINDEXES I,
```

```
      SYSIBM.SYSTABLES  T,
      SYSIBM.SYSTABLESPACE S
 WHERE I.CREATOR='DSN8610' AND
       I.NAME='XEMP1' AND
       I.TBCREATOR=T.CREATOR AND
       I.TBNAME=T.NAME AND
       T.DBNAME=S.DBNAME AND
       T.TSNAME=S.NAME;
```

Using DSN1COPY

This section describes the following tasks associated with running the DSN1COPY utility:

"Altering a table before running DSN1COPY"
"Checking for inconsistent data"
"Translating DB2 internal identifiers"
"Using an image copy as input to DSN1COPY" on page 583
"Resetting page log RBAs" on page 583
"Copying multiple data set table spaces" on page 583
"Restoring indexes with DSN1COPY" on page 583
"Restoring table spaces with DSN1COPY" on page 584
"Printing with DSN1COPY" on page 584
"Copying tables from one subsystem to another" on page 584

Altering a table before running DSN1COPY

Only the table's description changes when you use ALTER TABLE ADD COLUMN. You must run REORG on the table space (so the data matches its description) before you run DSN1COPY on the table space.

Checking for inconsistent data

When critical data is involved, use the CHECK option to prevent the undetected copying of inconsistent data to the output data set. The CHECK option performs validity checking on one page at a time.

You must run a CHECK utility job on the table space involved to ensure that no inconsistencies exist between data and indexes on that data:

- Before using DSN1COPY to save critical data that is indexed
- After using DSN1COPY to restore critical data that is indexed

The CHECK utility performs validity checking between pages.

Translating DB2 internal identifiers

If you use DSN1COPY to load data into a table space or index without specifying OBIDXLAT, be careful not to invalidate embedded DB2 internal identifiers. Those OBIDs can become invalid in the following ways:

- When you drop and recreate tables after you created the data saved by DSN1COPY but before you used it
- When a difference exists among the following attributes between the target subsystem and the source subsystem:
  - Table space attributes of BUFFERPOOL or NUMPARTS
  - Table attributes other than table name, table space name, and database name
  - The order of the table spaces, indexes, and tables the user defined or dropped in the source and target databases

To protect against invalidating the OBIDs, specify the OBIDXLAT parameter for DSN1COPY. The OBIDXLAT parameter translates OBID, DBID, or PSID before DSN1COPY copies the data.

Using an image copy as input to DSN1COPY

If you want to include FULLCOPY to use image copies as input to DSN1COPY, you must produce those image copies by using the COPY utility with SHRLEVEL REFERENCE. Using this parameter ensures that the data contained in your image copies is consistent. DSN1COPY accepts an index image copy as input when you specify the FULLCOPY option.

Resetting page log RBAs

The RESET option resets the log RBAs that are recorded in a table space or index space and the high formatted page number in the header page to 0. DSN1COPY performs CHECK processing whether or not you explicitly requested CHECK.

Do not specify the RESET parameter for page sets that are in group buffer pool RECOVER-pending (GRECP) status.

Copying multiple data set table spaces

When you use DSN1COPY to copy from an image copy of a table space's data sets to a table space's data sets, specify the following SYSUT2 data sets:

- **If SYSUT1 is an image of a single-partition**, SYSUT2 must list the data set name for that partition of the table space. Code the NUMPARTS(*nn*) parameter, where *nn* is the number of partitions in the whole table space.
- **If SYSUT1 is an image copy of a whole partitioned table space**, SYSUT2 must list the name of the table space's first data set. **Attention:** All data sets in the partitioned table space must use the same fifth level qualifier, I0001 or J0001, before DSN1COPY can run successfully on a partitioned table space. In this case, DSN1COPY allocates all of the target data sets. However, you must previously define the target data sets by using IDCAMS. Code the NUMPARTS(*nn*) parameter, where *nn* is the number of partitions in the whole table space.
- **If SYSUT1 is an image copy of a data set belonging to a multiple data set in a linear table space**, SYSUT2 should be the name of the output data set. Do not specify NUMPARTS, because this parameter is only for partitioned table spaces.
- **If SYSUT1 is an image copy of all data sets in a linear table space**, SYSUT2 should be the name of the table space's first data set. DSN1COPY allocates all target data sets.

Restoring indexes with DSN1COPY

When a table space has been restored using either the TOCOPY option of RECOVER or the DSN1COPY utility, you can restore the indexes in one of three ways:

- Use the RECOVER utility, if you have a full image copy available and the index was defined with the COPY YES option.
- An alternative is to use DSN1COPY on the indexes, if a copy is available. If you specified the OBIDXLAT option for the data, you must also specify the OBIDXLAT option for the indexes. Also, the user must have copied the indexes at the same time as the data; otherwise, inconsistencies may exist.

- If you do not have an image copy of the index, use the REBUILD INDEX utility, which will reconstruct the indexes from the data. However, for table spaces with millions of rows, REBUILD INDEX might take a long time. For more information about the REBUILD INDEX utility, refer to “Chapter 20. REBUILD INDEX” on page 259.

Restoring table spaces with DSN1COPY

It is impossible to use RECOVER TOCOPY for an image-copy data set that is not referenced by SYSIBM.SYSCOPY for that table space or data set. An attempt to do so results in the message `"TOCOPY DATASET NOT FOUND"`.

The MODIFY utility may have removed the row in SYSIBM.SYSCOPY. If this has happened, and the image copy is a full image copy with SHRLEVEL REFERENCE, DSN1COPY can restore the table space or data set.

DSN1COPY can restore the object to an incremental image copy, but it needs to have first restored the previous full image copy and any intermediate incremental image copies. This ensures data integrity. It is your responsibility to get the sequence of image copies right. DB2 cannot help ensure the proper sequence.

If you use DSN1COPY for point-in-time recovery, the table space becomes not recoverable. Because DSN1COPY executed outside of DB2’s control, DB2 is not aware that you recovered to a point-in-time. DSN1COPY recovers the affected table space after point-in-time recovery in conjunction with the following steps:

1. Remove old image copies by using MODIFY AGE.
2. Create one or more full image copies by using SHRLEVEL REFERENCE.

Printing with DSN1COPY

If you want to print one or more pages without having the copy function, use DSN1PRNT to avoid unnecessary reading of the input file.

When you use DSN1COPY for printing, you must specify the PRINT parameter. The requested operation takes place only for the data set specified. If the input data set belongs to a linear table space or index space that is larger than 2 gigabytes, you must specify the correct data set. Or, if it is a partitioned table space or partitioned index, you must specify the correct data set. For example, DSN1COPY prints a page range in the second partition of a four-partition table space. DSN1COPY does this by specifying NUMPARTS(4) and the data set name of the second data set in the VSAM group (DSN=...A002).

To print a full image copy data set (rather than recovering a table space), specify a DUMMY SYSUT2 DD card and specify the FULLCOPY parameter.

Copying tables from one subsystem to another

You must be careful when you copy a table containing an identity column from one DB2 subsystem to another:

1. Stop the table space on the source subsystem.
2. Issue a SELECT statement to query the SYSIBM.SYSSEQUENCES entry corresponding to the identity column for this table on the source subsystem. Add the INCREMENT value to the MAXASSIGNEDVAL to let DSN1COPY assign the next value (*nv*).
3. Create the table on the target subsystem. On the identity column specification, specify *nv* for the START WITH value, and ensure that all of the other identity column attributes are the same as for the source table.

4. Stop the table space on the target subsystem.
5. Copy the data by using DSN1COPY.
6. Start the table space on the source subsystem for read-write access.
7. Start the table space on the target subsystem for read-write access.

Sample control statements

If you run online REORG with the FASTSWITCH option, the fifth qualifier in the data set name can be either I0001 or J0001. Use I0001 in these examples.

Example 1: Running DSN1COPY with the CHECK option

```
//RUNCOPY EXEC PGM=DSN1COPY,PARM='CHECK'
//* COPY VSAM TO SEQUENTIAL AND CHECK PAGES
//STEPLIB  DD DSN=PDS CONTAINING DSN1COPY
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DSN=DSNCAT.DSNDBC.DSNDB01.SYSUTILX.I0001.A001,DISP=OLD
//SYSUT2   DD DSN=TAPE.DS,UNIT=TAPE,DISP=(NEW,KEEP),VOL=SER=UTLBAK
```

Example 2: Translation of the DB2 internal identifiers by using the OBIDXLAT parameter

```
//EXECUTE EXEC PGM=DSN1COPY,PARM='OBIDXLAT'
//STEPLIB  DD  DSN=PDS CONTAINING DSN1COPY
//SYSPRINT DD  SYSOUT=A
//SYSUT1   DD  DSN=DSNC710.DSNDBC.DSN8D71P.DSN8S71C.I0001.A001,
//             DISP=OLD
//SYSUT2   DD  DSN=DSNC618.DSNDBC.DSN8D71P.DSN8S71C.I0001.A001,
//             DISP=OLD
//SYSXLAT  DD  *
260,280
2,10
3,55
6,56
7,57
/*
```

Example 3: Printing a single page of a partitioned table space

```
//PRINT EXEC PGM=DSN1COPY,PARM='PRINT(2002A1),NUMPARTS(8)'
//* PRINT A PAGE IN THE THIRD PARTITION OF A TABLE SPACE CONSISTING
//* OF 8 PARTITIONS.
//SYSUDUMP DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSUT2   DD DUMMY
//SYSUT1   DD DSN=DSNCAT.DSNDBD.MMRDB.PARTEMP1.I0001.A003,DISP=OLD
```

Example 4: To print 16 pages of a nonpartitioning index

```
//PRINT2 EXEC PGM=DSN1COPY,PARM=(PRINT(F0000,F000F),PIECESIZ(64M))
//* PRINT 16 PAGES IN THE 61ST PIECE OF AN NPI WITH PIECE SIZE OF 64M
//SYSUDUMP DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSUT2    DD DUMMY
//SYSUT1    DD DISP=OLD,DSN=DSNCAT.DSTDBD.MMRDB.NPI1.I0001.A061
```

Example 5: Copying individual partitions of a partitioned table space

In this example, partition 3 and partition 4 are copied from image copy data sets into the partitioned table space. The fifth qualifier in the data set names can be different because each job step lists an individual partition.

```
//STEP1    EXEC PGM=DSN1COPY,
//            PARM='NUMPARTS(16),RESET,FULLCOPY'
//SYSUDUMP DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DISP=SHR,DSN=PROD.IMAGE.COPY.PART3
//SYSUT2   DD DISP=OLD,DSN=DSNCAT.DSNDBD.TESTDB.TS01.I0001.A003
//STEP2    EXEC PGM=DSN1COPY,
//            PARM='NUMPARTS(16),RESET,FULLCOPY'
//SYSUDUMP DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DISP=SHR,DSN=PROD.IMAGE.COPY.PART4
//SYSUT2   DD DISP=OLD,DSN=DSNCAT.DSNDBD.TESTDB.TS01.J0001.A004
```

Example 6: Copying all partitions of a partitioned table space

This example copies data into all partitions of a partitioned table space using an image copy of the entire table space as input. You must ensure that the fifth qualifier in the data set name is the same, either I0001 or J0001 for **all** partitions of the output table space before running this type of job stream.

```
//DSN1COPY EXEC PGM=DSN1COPY,
//            PARM='NUMPARTS(16),RESET,FULLCOPY'
//SYSUDUMP DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DISP=SHR,DSN=PROD.IMAGE.COPY.DSNUMALL
//SYSUT2   DD DISP=OLD,DSN=DSNCAT.DSNDBD.TESTDB.TS01.I0001.A001
```

DSN1COPY output

Interpreting output

One intended use of this utility is to aid in determining and correcting system problems. When diagnosing DB2, you might need to refer to licensed documentation to interpret output from this utility.

Chapter 37. DSN1LOGP

The DSN1LOGP utility formats the contents of the recovery log for display. The two recovery log report formats are:

- A *detail report* of individual log records. This information helps IBM Support Center personnel analyze the log in detail. (This book does not include a full description of the detail report.)
- A *summary report* helps you:
  - Perform a conditional restart
  - Resolve indoubt threads with a remote site
  - Detect problems with data propagation

You can specify the range of the log to process and select criteria within the range to limit the records in the detail report. For example, you can specify:

- One or more units of recovery identified by URID
- A single database

By specifying a URID and a database, you can display recovery log records that correspond to the use of one database by a single unit of recovery.

Syntax and options of the control statement

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

© Copyright IBM Corp. 1983, 2001

DSN1LOGP syntax diagram

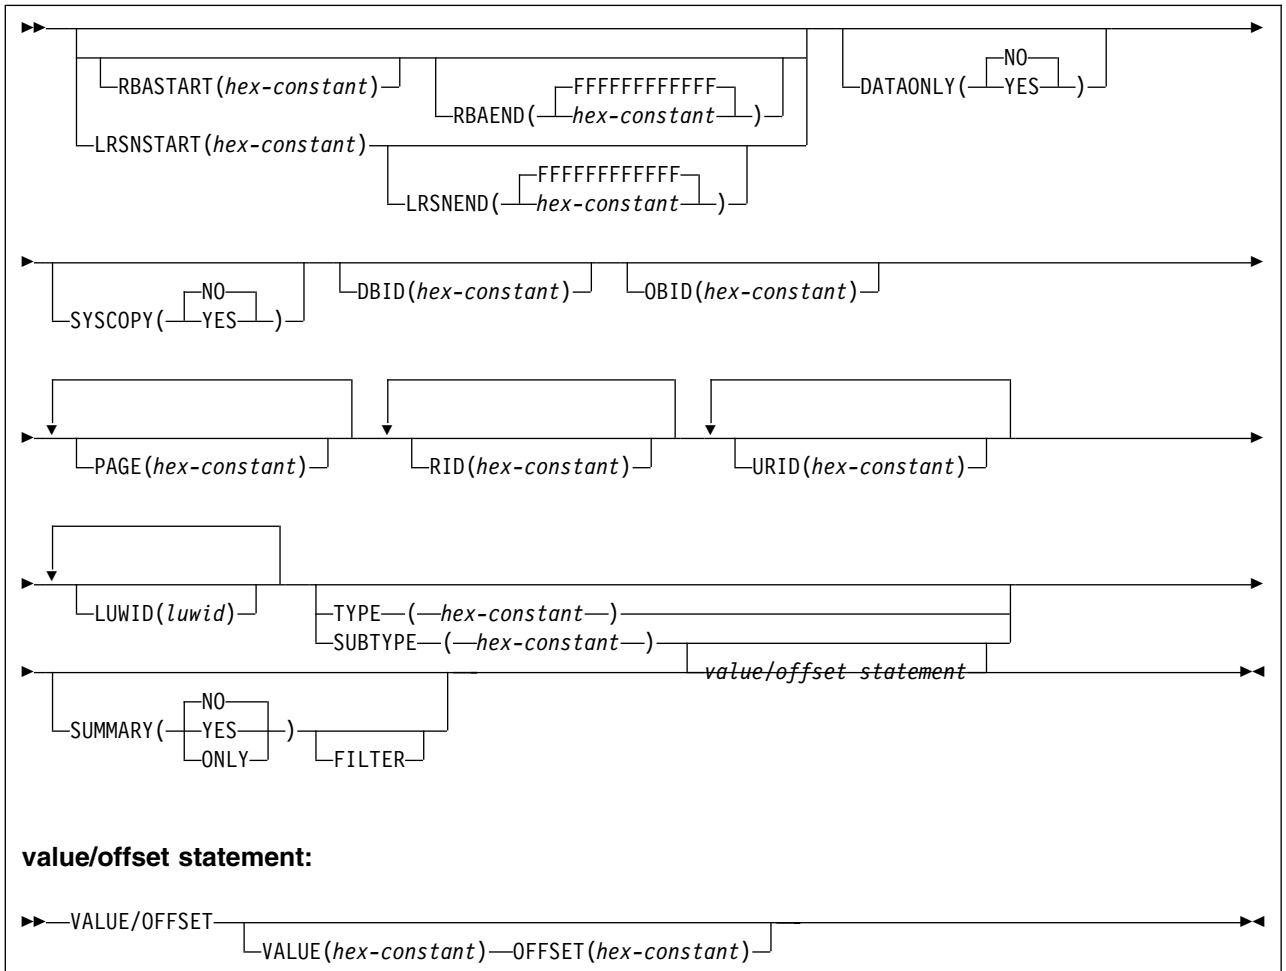

Attention: If you specify more than one parameter:

- Separate them by commas (no blanks)
- Specify them in any order.

Option descriptions

To execute DSN1LOGP, construct a batch job. The utility name, DSN1LOGP, should appear on the EXEC statement, as shown in “Sample control statements” on page 596.

Specify keywords in up to 50 control statements in the SYSIN file. Each control statement can have up to 72 characters. To specify no keywords, either use a SYSIN file with no keywords following it or omit the SYSIN file from the job JCL.

The keywords are described below; alternative spellings or abbreviations are noted.

You can include blanks between keywords, and also between the keywords and the corresponding values.

RBASTART(*hex-constant***)**
Specifies the hexadecimal log RBA from which to begin reading. If the value does not match the beginning RBA of one of the log records, DSN1LOGP begins reading at the beginning RBA of the next record. For any given job, specify this keyword only once. Alternative spellings: STARTRBA, ST.

hex-constant is a hexadecimal value consisting of 1 to 12 characters (6 bytes), and leading zeros are not required.

The **default** is **0**.

RBAEND(*hex-constant***)**
Specifies the last valid hexadecimal log RBA to extract. If the specified RBA is in the middle of a log record, DSN1LOGP continues reading the log in an attempt to return a complete log record.

To read to the last valid RBA in the log, specify RBAEND(FFFFFFFFFFFF). For any given job, specify this keyword only once. Alternative spellings: ENDRBA, EN.

hex-constant is a hexadecimal value consisting of 1 to 12 characters (6 bytes), and leading zeros are not required.

The **default** is **FFFFFFFFFFFF**.

RBAEND can be specified only if RBASTART is specified.

LRSNSTART(*hex-constant***)**
Specifies the log record sequence number (LRSN) from which to begin the log scan. DSN1LOGP starts its processing on the first log record containing an LRSN value greater than or equal to the LRSN value specified on LRSNSTART. The default LRSN is the LRSN at the beginning of the data sets. Alternative spellings: STARTLRSN, STRTLRSN, and LRSNSTRT.

For any given job, specify this keyword only once.

You must specify this keyword to search the member BSDSs and locate the log data sets from more than one DB2 subsystem. You can specify either the LRSNSTART keyword or the RBASTART keyword to search the BSDS of a single DB2 subsystem and locate the log data sets.

LRSNEND(*hex-constant***)**
Specifies the LRSN value of the last log record to be scanned. When LRSNSTART is specified, the default is X'FFFFFFFFFFFF'. Otherwise, it is the end of the data sets. Alternative spelling: ENDLRSN.

For any given job, specify this keyword only once.

DATAONLY Limits the log records in the detail report to those that represent data changes (insert, page repair, update space map, and so on).

The **default** is **DATAONLY**(**NO**).

(YES) Extracts log records for data changes only. For example, DATAONLY(YES), together with a DBID and OBID, reads only the log records that modified data for that DBID and OBID.

(NO) Extracts all record types.

SYSCOPY Limits the detail report to SYSCOPY log records.

The **default** is **SYSCOPY**(**NO**).

(YES) Includes only SYSCOPY log records in the detail report.

(NO) Does not limit records to SYSCOPY records only.

DBID(*hex-constant***)**

Specifies a hexadecimal database identifier (DBID). DSN1LOGP extracts only records associated with that DBID. For any given job, specify this keyword only once.

hex-constant is a hexadecimal value consisting of 1 to 4 characters. Leading zeros are not required.

The DBID is displayed in many DB2 messages. You can also find the DBID in the DB2 catalog for a specific object (for example, in the column named "DBID" of the SYSIBM.SYSTABLESPACE catalog table).

When you select a DBID from a catalog table, the value is displayed in decimal format. Use the following SQL HEX function in a SELECT statement to convert a DBID to hexadecimal format:

```
SELECT NAME, DBNAME, HEX(DBID), HEX(PSID)
FROM SYSIBM.SYSTABLESPACE
WHERE NAME ='table space name'

SELECT NAME, DBNAME, HEX(DBID), HEX(ISOBID)
FROM SYSIBM.SYSINDEXES
WHERE NAME ='index name'
```

OBID(*hex-constant***)**

Specifies a hexadecimal database object identifier, either a data page set identifier (PSID) or an index page set identifier (ISOBID). DSN1LOGP extracts only records associated with that identifier.

hex-constant is a hexadecimal value consisting of 1 to 4 characters. Leading zeros are not required.

Whenever DB2 makes a change to data, the log record describing the change identifies the database by DBID and the table space by page set ID (PSID). You can find the PSID column in the SYSIBM.SYSTABLESPACE catalog table.

You can also find a column named OBID in the SYSIBM.SYSTABLESPACE catalog table. That column actually contains the OBID of a file descriptor; don't confuse this with the PSID, which is the information you must include when you execute DSN1LOGP.

Whenever DB2 makes a change to an index, the log record describing the change identifies the database (by DBID) and the index space (by index space OBID, or ISOBID). You can find the ISOBID for an index space in the column named ISOBID in the SYSIBM.SYSINDEXES catalog table.

You will also find a column named OBID in the SYSIBM.SYSINDEXES catalog table. This column actually contains the identifier of a fan set descriptor; don't confuse this with the ISOBID, which is the information you must include when you execute DSN1LOGP.

When you select either the PSID or the ISOBID from a catalog table, the value is displayed in decimal format. Use the SQL HEX function in your select statement to convert them to hexadecimal.

For any given DSN1LOGP job, use this keyword only once. If you specify OBID, you must also specify DBID.

PAGE(*hex-constant***)**
Specifies a hexadecimal page number. When data or an index is changed, a recovery log record is written to the log, identifying the object identifier and the page number of the changed data page or index page. Specifying a page number limits the search to a single page; otherwise, all pages for a given combination of DBID and OBID are extracted. The log output also contains page set control log records for the specified DBID and OBID, as well as the system event log records, unless DATAONLY(YES) is also specified.

hex-constant is a hexadecimal value consisting of a maximum of eight characters.

You can specify a maximum of 100 PAGE keywords in any given DSN1LOGP job. You must also specify the DBID and OBID keywords that correspond to those pages.

The PAGE and RID keywords cannot both be specified.

RID(*hex-constant***)**
Specifies a record identifier, which is a hexadecimal value consisting of 10 characters, with the first eight characters representing the page number and the last two characters representing the page ID map entry number. The option limits the log records extracted to those associated with that particular record. The log records extracted include not only those directly associated with the RID, such as insert and delete, but also the control records associated with the DBID and OBID specifications, such as page set open, page set close, set write, reset write, page set write, data set open, and data set close.

You can specify a maximum of 40 RID keywords in any given DSN1LOGP job. You must also specify the DBID and OBID keywords that correspond to the specified records.

The PAGE and RID keywords are mutually exclusive.

URID(*hex-constant***)**
Specifies a hexadecimal unit of recovery identifier (URID). Changes to data and indexes occur in the context of a DB2 unit of recovery, which is identified on the log by a BEGIN UR record. In the summary DSN1LOGP report , the URID is listed in the STARTRBA field in message DSN1162I. In the detail DSN1LOGP report, look under the subtype of BEGIN UR and the URID is listed in the URID field. Using the log RBA of that record as the URID value limits the extraction of information from the DB2 log to that unit of recovery.

hex-constant is a hexadecimal value consisting of 1 to 12 characters (6 bytes). Leading zeros are not required.

You can specify a maximum of 10 URID keywords in any given DSN1LOGP job.

LUWID(*luwid***)** Specifies up to 10 LUWIDs to include information about in the summary report.

luwid consists of three parts: an LU network name, an LUW instance number, and a commit sequence number. If you supply the first two parts, the summary report includes an entry for each commit performed in the logical unit of work (within the search range). If you supply all three parts, the summary report includes an entry for only that LUWID.

The LU network name consists of a one- to eight-character network ID, a period, and a one- to eight-character network LU name. The LUW instance number consists of a period followed by 12 hex characters. The last element of the LUWID is the commit sequence number of 4 hex characters, preceded by a period.

TYPE(*hex-constant***)**

Limits the log records extracted to records of a specified type. The TYPE and SUBTYPE options are mutually exclusive.

hex-constant indicates the type, as follows:

Constant	Description
2	Page set control record
4	SYSCOPY utility record
10	System event record
20	UR control record
100	Checkpoint record
200	UR-UNDO record
400	UR-REDO record
800	Archive quiesce record
1000 to 8000	Assigned by the resource manager

SUBTYPE(*hex-constant***)**

Restricts formatting to a particular subtype of unit of recovery undo and redo log records (types 200 and 400). The TYPE and SUBTYPE options are mutually exclusive.

hex-constant indicates the subtype, as follows:

Constant	Description
1	Update data page
2	Format page or update space map
3	Update space map bits
4	Update to index space map
5	Update to index page
6	DBA table update log record
7	Checkpoint DBA table log record
9	DBD virtual memory copy
A	Exclusive lock on page set partition or DBD
B	Format file page set
C	Format index page set
F	Update by repair (first half if 32 KB)
10	Update by repair (second half if 32 KB)
11	Allocating or deallocating a segment entry
12	Undo/redo log record for modified page or redo log record for formatted page
14	Savepoint
15	Other DB2 component log records written for RMID 14
17	Checkpoint record of modified page set
19	Type 2 Index update

1A Type 2 Index under/redo or redo log record
1B Type 2 Index change notification log record
1C Type 2 Index space map update
1D DBET log record with exception data
1E DBET log record with LPL/GRECP data
65 Data propagation diagnostic log
81 Index dummy compensation log record
82 START DATABASE ACCESS (FORCE) log record

The VALUE and OFFSET options must be used together. You can specify a maximum of 10 VALUE/OFFSET pairs. The SUBTYPE parameter is required when using the VALUE and OFFSET options.

VALUE(*hex-constant***)**
Specifies a value that must appear in a log record to be extracted.

hex-constant is a hexadecimal value consisting of a maximum of 64 characters and must be an even number of characters.

The SUBTYPE keyword must be specified before the VALUE option.

OFFSET(*hex-constant***)**
Specifies an offset from the log record header at which the value specified in the VALUE option must appear.

hex-constant is a hexadecimal value consisting of a maximum of 8 characters.

The SUBTYPE keyword must be specified before specifying the OFFSET option.

SUMMARY Summarizes all recovery information within the RBASTART and RBAEND specifications. You can use summary information to determine what work is incomplete when DB2 starts. You cannot limit the output of the summary report with any of the other options, except by using the FILTER option with a URID or LUWID specification.

The **default** is **SUMMARY**(**NO**).

(YES) Generates both a detail and summary report.

(NO) Generates only a detail report.

(ONLY)
Generates only a summary report.

FILTER Restricts the summary report to include messages for only the specified URIDs and LUWIDs. Specify this option only once.

The SUMMARY keyword must be specified before FILTER.

Before running DSN1LOGP

This section contains information needed before running DSN1LOGP.

Environment

DSN1LOGP runs as a batch Multiple Virtual Storage (MVS) job.

DSN1LOGP runs on archive data sets, but not active data sets, when DB2 is running.

Authorization required

DSN1COPY does not require authorization. However, any data sets protected by Resource Access Control Facility (RACF) require RACF authorization.

Control statement

See “Syntax and options of the control statement” on page 587 for DSN1LOGP syntax and option descriptions.

Required data sets

When you execute DSN1LOGP, you must provide the following data definition (DD) statements:

SYSPRINT DSN1LOGP writes all error messages, exception conditions, and the detail report to the SYSPRINT file. The logical record length (LRECL) is 131.

SYSIN DSN1LOGP specifies keywords in this file. DSN1LOGP describes the control statement keywords under “Option descriptions” on page 588. The LRECL must be 80. Keywords and values must appear in characters 1 through 72. DSN1LOGP allows specification of as many as 50 control statements for a given job. DSN1LOGP concatenates all records into a single string.

SYSSUMRY DSN1LOGP writes the formatted output of a summary report to the SYSSUMRY file. The LRECL is 131. For an example of the appropriate job control language (JCL), see page 598.

DSN1LOGP identifies the recovery log by DD statements that are described in the stand-alone log services. For a description of these services, see Appendix C (Volume 2) of *DB2 Administration Guide*.

Identifying log data sets

You must identify to DSN1LOGP the log data sets to process by including at least one of the following DD statements.

BSDS The BSDS identifies and provides information about all active log data sets and archive log data sets that exist in your DB2 subsystem. When you identify the BSDS to DSN1LOGP, you must provide the beginning and ending relative byte addresses for the range of the recovery log you want displayed. DSN1LOGP then associates the beginning RBA specifications and the ending RBA specifications with the appropriate data set names.

See Example 1 on page 596 for guidance in using this DD statement.

ACTIVE*n* If the BSDS is not available, and if the active log data sets involved have been copied and sent to you, you can specify the set of active log data sets to be processed by DSN1LOGP by specifying one or more ACTIVE DD statements. If you used the REPRO command of access method services for copying the active log, you must identify this data set in an ARCHIVE DD statement.

Each DD statement that you include identifies another active log data set. If you identify more than one active log data set, you must list the ACTIVE*n* DD statements in ascending log RBA sequence.

For example, ACTIVE1 must identify a portion of the log that is less than ACTIVE2. And, ACTIVE2 must identify a portion of the log that is less than ACTIVE3. If you do not specify this correctly, errors that DSN1LOGP does not detect can occur.

When you identify active log data sets, you do not need to use the RBASTART and RBAEND keywords (as you do when you identify the BSDS). DSN1LOGP will scan all active log data sets that the job indicates only when the data sets are in the correct log RBA sequence.

See Example 2 on page 597 for guidance in using these DD statements.

ARCHIVE If the BSDS is not available (as described under ACTIVE*n*, above), you can specify which archive log data sets are to be processed by specifying one ARCHIVE DD statement, concatenated with one or more DD statements as shown in Example 3 on page 597.

Each DD statement you include identifies another archive log data set. If you identify more than one archive log data set, you must list the DD statements corresponding to the multiple archive log data sets in ascending log RBA sequence. If you do not specify this correctly, errors that DSN1LOGP does not detect can occur.

When you identify archive log data sets, you do not need to use the RBASTART and RBAEND keywords. DSN1LOGP scans all archive log data sets indicated by the job only when the data sets are in the correct log RBA sequence.

See Example 3 on page 597 for guidance in using the ARCHIVE DD statement.

Data sharing requirements: When selecting log records from more than one DB2 subsystem, you must use all of the following DD statements to locate the log data sets:

GROUP
M*xx*BSDS
M*xx*ARCHV
M*xx*ACT*n*

See Appendix C (Volume 2) of *DB2 Administration Guide* for descriptions of those statements. If you use GROUP or M*xx*BSDSs to locate the log data sets, you must use LRSNSTART to define the selection range.

Using DSN1LOGP

This section describes the following tasks associated with running the DSN1LOGP utility:

"Reading archive log data sets on tape"
"Locating table and index identifiers" on page 596

Reading archive log data sets on tape

If you stored your archive logs on tape, DSN1LOGP constructed two files on tape during the archiving process. The first file is the BSDS, and the second is a dump of the active log that DSN1LOGP is currently archiving. If a failure occurs during the time DSN1LOGP is archiving the BSDS, DB2 might omit the BSDS. In this case, the first file contains the active log.

If you perform archiving on tape, the first letter of the lowest-level qualifier varies for both the first and second data sets. The first letter of the first data set is B (for BSDS), and the first letter of the second data set is A (for archive). Hence, the data set names all end in A*xxxxxxx*, and the DD statement identifies each of them as the second data set on the corresponding tape:

```
LABEL=(2,SL)
```

When reading archive log data sets on tape (or copies of active log data sets on tape), add one or more of the following Job Entry Subsystem (JES) statements:

JES2 environment JCL	Description
/*SETUP	Alert the MVS operator to prepare to mount a specified list of tapes.
/*HOLD	Place the job in HOLD status until the operator has located the tapes and is ready to release the job.
TYPRUN=HOLD	Perform the same function as /*HOLD. The system places the JCL on the JOB card.

JES3 Environment JCL	Description
//*MAIN SETUP=JOB	Alert the MVS operator to mount the initial volumes before the job executes.
//*MAIN HOLD=YES	Place the job in HOLD status until the operator is ready to release the job.
TYPRUN=HOLD	Perform the same function as //*MAIN HOLD=YES. The system places the JCL on the JOB card.

Alternatively, the user submits the job to an MVS initiator that your operations center has established for exclusive use by jobs that require tape mounts. The user specifies the initiator class using the CLASS parameter on the JOB card, in both JES2 and JES3 environments.

For additional information on these options, refer to the *OS/390 MVS JCL User's Guide* or the *OS/390 MVS JCL Reference.*

Locating table and index identifiers

You can use the DSN1PRNT utility to find the DBIDs, PSIDs, ISOBIDs, and OBIDs of the tables and indexes from the system tables. For more information, see “Chapter 38. DSN1PRNT” on page 609.

Sample control statements

Example 1: Using DSN1LOGP with an available BSDS. This example shows how to extract the information from the recovery log when you have the BSDS available. The extraction starts at the log RBA of X'AF000' and ends at the log RBA of X'B3000'. DSN1LOGP identified the table or index space by the DBID of X'10A' (266 decimal) and the OBID of X'1F' (31 decimal).

```
//STEP1 EXEC PGM=DSN1LOGP
//STEPLIB  DD DSN=PDS containing DSN1LOGP
//SYSPRINT DD SYSOUT=A
//SYSABEND DD SYSOUT=A
//BSDS     DD DSN=DSNCAT.BSDS01,DISP=SHR
```

```
//SYSIN    DD *
  RBASTART (AF000) RBAEND (B3000)
  DBID (10A) OBID(1F)
/*
```

You can think of the DB2 recovery log as a large sequential file. Whenever recovery log records are written, they are written to the end of the log. A log RBA is the address of a byte on the log. Because the recovery log is larger than a single data set, the recovery log is physically stored on many data sets. DB2 records the RBA ranges and their corresponding data sets in the BSDS. To determine which data set contains a specific RBA, read the information about the DSNJU004 utility on page 547 and see Part 4 (Volume 1) of *DB2 Administration Guide*. During normal DB2 operation, messages are issued that include information about log RBAs.

Example 2: Using DSN1LOGP on the active log (no BSDS available). This example shows how to extract the information from the active log when the BSDS is not available. The extraction includes log records that apply to the table space or index space identified by the DBID of X'10A' and the OBID of X'1F'. The only information that is extracted is information relating to page numbers X'3B' and X'8C'. You can omit beginning and ending RBA values for ACTIVE*n* or ARCHIVE DD statements, because the DSN1LOGP search includes all specified ACTIVE*n* DD statements. The DD statements ACTIVE1, ACTIVE2, and ACTIVE3 specify the log data sets in ascending log RBA range. Use the DSNJU004 utility to determine what the log RBA range is for each active log data set. If the BSDS is not available and you cannot determine the ascending log RBA order of the data sets, you must run each log data set individually.

```
//STEP1 EXEC PGM=DSN1LOGP
//STEPLIB  DD DSN=PDS containing DSN1LOGP
//SYSPRINT DD SYSOUT=A
//SYSABEND DD SYSOUT=A
//ACTIVE1  DD DSN=DSNCAT.LOGCOPY1.DS02,DISP=SHR      RBA X'A000' - X'BFFF'
//ACTIVE2  DD DSN=DSNCAT.LOGCOPY1.DS03,DISP=SHR      RBA X'C000' - X'EFFF'
//ACTIVE3  DD DSN=DSNCAT.LOGCOPY1.DS01,DISP=SHR      RBA X'F000' - X'12FFF'
//SYSIN    DD *
  DBID (10A) OBID(1F) PAGE(3B) PAGE(8C)
/*
```

Example 3: Using DSN1LOGP on archive log data (no BSDS available). This example shows how to extract the information from archive logs when the BSDS is not available. The extraction includes log records that apply to a single unit of recovery (whose URID is X'61F321'). Because the BEGIN UR is the first record for the unit of recovery and is at X'61F321', the beginning RBA is specified to indicate that it is the first RBA in the range from which to extract recovery log records. Also, because no ending RBA value is specified, all specified archive logs are scanned for qualifying log records. The specification of DBID(4) limits the scan to changes that the specified unit of recovery made to all table spaces and index spaces in the database whose DBID is X'4'.

```
//STEP1 EXEC PGM=DSN1LOGP
//STEPLIB  DD DSN=PDS containing DSN1LOGP
//SYSPRINT DD SYSOUT=A
//SYSABEND DD SYSOUT=A
//ARCHIVE  DD DSN=DSNCAT.ARCHLOG1.A0000037,UNIT=TAPE,VOL=SER=T10067,
//            DISP=(OLD,KEEP),LABEL=(2,SL)
//         DD DSN=DSNCAT.ARCHLOG1.A0000039,UNIT=TAPE,VOL=SER=T30897,
//            DISP=(OLD,KEEP),LABEL=(2,SL)
//         DD DSN=DSNCAT.ARCHLOG1.A0000041,UNIT=TAPE,VOL=SER=T06573,
//            DISP=(OLD,KEEP),LABEL=(2,SL)
```

```
//SYSIN    DD *
  RBASTART (61F321)
  URID (61F321) DBID(4)
/*
```

Example 4: Using DSN1LOGP with the SUMMARY option. The DSN1LOGP SUMMARY option allows you to scan the recovery log to determine what work is incomplete at restart time. You can specify this option either by itself or when you use DSN1LOGP to produce a detail report of log data. Summary log results appear in SYSSUMRY; therefore, you must include a SYSSUMRY DD statement as part of the JCL with which you execute DSN1LOGP.

This example produces both a detail and a summary report using the BSDS to identify the log data sets. The summary report summarizes all recovery log information within the RBASTART and RBAEND specifications. You cannot limit the output of the summary report with any of the other options, except by using the FILTER option with a URID or LUWID specification. RBASTART and RBAEND specification use depends on whether a BSDS is used.

This example is similar to Example 1, in that it shows how to extract the information from the recovery log when you have the BSDS available. However, this example also shows you how to specify a summary report of all logged information between the log RBA of X'AF000' and the log RBA of X'B3000'. This summary is generated with a detail report, but will be printed to SYSSUMRY separately.

```
//STEP1 EXEC PGM=DSN1LOGP
//STEPLIB  DD DSN=PDS containing DSN1LOGP
//SYSPRINT DD SYSOUT=A
//SYSSUMRY DD SYSOUT=A
//SYSABEND DD SYSOUT=A
//BSDS     DD DSN=DSNCAT.BSDS01,DISP=SHR
//SYSIN    DD *
  RBASTART (AF000) RBAEND (B3000)
  DBID (10A) OBID(1F) SUMMARY(YES)
/*
```

Example 5: Data sharing— using DSN1LOGP on all members of a data sharing group. This example shows extract log information pertaining to the table space identified by DBID X'112' and OBID X'1D', from all members of a data sharing group.

```
//STEP1 EXEC PGM=DSN1LOGP
//STEPLIB DD  DSN=PDS containing DSN1LOGP
//SYSPRINT SYSOUT=A
//SYSABEND SYSOUT=A
//GROUP DD DSN=DSNDB0G.BSDS01,DISP=SHR
//SYSIN DD *
  DATAONLY (YES)
  LRSNSTART (A7951A001AD5) LRSNEND (A7951A003B6A)
  DBID (112) OBID(1D)
/*
```

Example 6: Data sharing— using DSN1LOGP on a single member of a data sharing group. This example shows extract log information pertaining to the table space identified by DBID X'112' and OBID X'1D', from a single member of a data sharing group.

```
//STEP1 EXEC PGM=DSN1LOGP
//STEPLIB DD  DSN=PDS containing DSN1LOGP
//SYSPRINT SYSOUT=A
//SYSABEND SYSOUT=A
//M01BSDS DD DSN=DSNDB0G.DB1G.BSDS01,DISP=SHR
//SYSIN DD *
```

```
  DATAONLY (YES)
  LRSNSTART (A7951A001AD5) LRSNEND (A7951A003B6A)
  DBID (112) OBID(1D)
/*
```

DSN1LOGP output

Interpreting output

One intended use of this utility is to aid in determining and correcting system problems. When diagnosing DB2, you might need to refer to licensed documentation to interpret output from this utility.

Reviewing DSN1LOGP output

With the SUMMARY option, you can produce a summary report, a detail report, or both.

Figure 40 on page 601 shows a sample of the summary report. Figure 41 on page 602 shows a sample of the detail report. Figure 42 on page 607 shows a sample of data propagation information from a summary report. A description of the output precedes each sample.

Description of the summary report

The summary report on page Figure 40 on page 601 contains a summary of completed events, consisting of an entry for each completed unit of work. Each entry shows, among other information, the start time, user, and all page sets that were modified.

The summary report is divided into two distinct sections:

- The first section is headed by the message:

  `DSN1150I SUMMARY OF COMPLETED EVENTS`
- The second section is headed by the message:

  `DSN1157I RESTART SUMMARY`

The first section lists all completed units of recovery (URs) and checkpoints within the range of the log scanned. Events are listed chronologically, with URs listed by the order of their completion and checkpoints listed when the end of a checkpoint is processed. The page sets changed by each completed UR are listed. If a log record associated with an UR is unavailable (as it would be if, for example, the range of the log scanned is not large enough to contain all records for a given UR), the attribute INFO=PARTIAL is displayed for the UR. Otherwise, the UR is marked INFO=COMPLETE.

The DISP attribute can be one of the following: COMMITTED, ABORTED, INFLIGHT, IN-COMMIT, IN-ABORT, POSTPONED ABORT, or CANCELLED. The DISP attributes COMMITTED and ABORTED are used in the first section; the remaining attributes are used in the second section.

The list in the second section shows the work required of DB2 at restart as it is recorded in the log you specified. If the log is available, the checkpoint to be used is identified, as is each outstanding UR together with the page sets it changed. Each page set with pending writes is also identified, as is the earliest log record required to complete those writes. If a log record associated with an UR is

unavailable, the attribute INFO=PARTIAL is displayed, and the identification of modified page sets is incomplete for that UR.

```
DSN1212I DSN1LGRD FIRST LOG LRSN ENCOUNTERED AA526968220D

================================================================================
DSN1150I SUMMARY OF COMPLETED EVENTS

DSN1151I DSN1LPRT  MEMBER=V61B     UR  CONNID=V61B      CORRID=021.OPNLGR00  AUTHID=SYSOPR    PLAN=SYSTEM
             START DATE=94.347 TIME=11:15:22  DISP=COMMITTED   INFO=COMPLETE
             STARTRBA=00000000E570  ENDRBA=00000000EB64  STARTLRSN=AA52696B1269  ENDLRSN=AA526999D14D  NID=*
             LUWID=USIBMSY.SYEC1B.AA52696825CE.0001  COORDINATOR=*
             PARTICIPANTS=*
             DATA MODIFIED:
                  DATABASE=0001=DSNDB01   PAGE SET=00CF=SYSLGRNX
                  DATABASE=0001=DSNDB01   PAGE SET=0087=DSNLLX01
                  DATABASE=0001=DSNDB01   PAGE SET=0086=DSNLLX02

DSN1151I DSN1LPRT  MEMBER=V61B     UR  CONNID=V61B      CORRID=021.OPNLGR00  AUTHID=SYSOPR    PLAN=SYSTEM
             START DATE=94.347 TIME=11:16:14  DISP=COMMITTED   INFO=COMPLETE
             STARTRBA=00000000ECFC  ENDRBA=00000000F20A  STARTLRSN=AA52699C97A9  ENDLRSN=AA52699CADC5  NID=*
             LUWID=USIBMSY.SYEC1B.AA52699C9508.0001  COORDINATOR=*
             PARTICIPANTS=*
             DATA MODIFIED:
                  DATABASE=0001=DSNDB01   PAGE SET=00CF=SYSLGRNX
                  DATABASE=0001=DSNDB01   PAGE SET=0087=DSNLLX01
                  DATABASE=0001=DSNDB01   PAGE SET=0086=DSNLLX02
....

DSN1213I DSN1LGRD LAST LOG LRSN ENCOUNTERED AA527C9B8392

DSN1214I NUMBER OF LOG RECORDS READ 0000000000004991

================================================================================
DSN1157I RESTART SUMMARY

DSN1153I DSN1LSIT CHECKPOINT  MEMBER=V61B
               STARTRBA=000000068CD3  ENDRBA=00000006CAED  STARTLRSN=AA527AA809DF  ENDLRSN=AA527AA829F4
               DATE=94.347  TIME=12:32:29

DSN1162I DSN1LPRT  MEMBER=V61C     UR  CONNID=BATCH     CORRID=S5529927      AUTHID=ADMF001   PLAN=PLNFW543
             START DATE=94.347 TIME=12:41:04  DISP=INFLIGHT    INFO=COMPLETE
             STARTRBA=000000016000  STARTLRSN=AA527C9278DF  NID=*
             LUWID=USIBMSY.SYEC1C.AA527C22E283.0001  COORDINATOR=*
             PARTICIPANTS=*
             DATA MODIFIED:
                  DATABASE=0113=DBFW5401  PAGE SET=0002=TPFW5401
                  DATABASE=0113=DBFW5401  PAGE SET=0005=IPFW5401

DSN1162I DSN1LPRT  MEMBER=V61A     UR  CONNID=BATCH     CORRID=S5529925      AUTHID=ADMF001   PLAN=PLNFW541
             START DATE=94.347 TIME=12:41:04  DISP=INFLIGHT    INFO=COMPLETE
             STARTRBA=000001F9A3C1  STARTLRSN=AA527C92E419  NID=*
             LUWID=USIBMSY.SYEC1DB2.AA527C1D674B.0001  COORDINATOR=*
             PARTICIPANTS=*
             DATA MODIFIED:
                  DATABASE=0113=DBFW5401  PAGE SET=0002=TPFW5401
 ...
 DSN1160I DATABASE WRITES PENDING:
                  DATABASE=0001=DSNDB01    PAGE SET=0046=DSNLUX02   START=000000068CD3
                  DATABASE=0001=DSNDB01    PAGE SET=0044=DSNLUX01   START=000000068CD3
 ...
                  DATABASE=0006=DSNDB06    PAGE SET=0076=DSNUCX01   START=000000068CD3
                  DATABASE=0006=DSNDB06    PAGE SET=0072=DSNUCH01   START=000000068CD3
 ...
```

Figure 40. Sample DSN1LOGP summary report

Description of the detail report

The *detail report* on page Figure 41 includes the following records:

- *Redo/undo log records*
- *System events log records*, including begin and end checkpoint records, begin current status rebuild records, and begin forward and backward recovery records
- *Page set control log records*, including open and close page set log records, open and close data set log records, set write, reset write, and page set write log records
- *UR control log records for the complete or incomplete unit of recovery*

You can reduce the volume of the detail log records by specifying an optional parameter.

```
DSN1212I DSN1LGRD FIRST LOG  RBA ENCOUNTERED 00000335916E

0000033591D4  MEMBER(M01     )  LRSN(AB62536BE583)  DBID(0006)  OBID(00B2)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET STATUS RECORD)
     *LRH* 00660066 00020009 0E800000 00000000 00000335 916E0126 00000335 916EAB62
           536BE583 0001
     0000  000600B2 C4E2D5C4 C2F0F640 C4E2D5E3 D5E7F0F1 00010000 92018000 00000334
     0020  EC3AAB62 5260AB0B 00000000 00000000 00000000 00000000 00000000 00000000

0000000109E2  MEMBER(M02     )  LRSN(AB6253746CE3)  DBID(0113)  OBID(0008)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET OPEN)
     *LRH* 00A0006E 00020001 0E800000 00000000 00000000 00000126 00000000 0000AB62
           53746CE3 0002
     0000  01130008 6C010100 00000005 0040C4C2 C6E6F0F0 F1F1C9C3 C6E6F0F0 F0F10001
     0020  00060000 10009201 00130020 00000000 00000000 00000000 00000000 00000000
     0040  00000000 00000010 00010000 00000000 00000000 00000000 00000000 00000000
     0060  00AB624B 192CEEAB 624B4783 F8000000 0000C4E2 D5C3F4F1 F040

000000010A82  MEMBER(M02     )  URID(000000010A82)  LRSN(AB6253747801)
     TYPE(UR CONTROL)  SUBTYPE(BEGIN UR)
     *LRH* 009000A0 00200001 03800000 00010A82 00000000 00000126 00000000 0000AB62
           53747801 0002
     0000  00010000 0000D000 00000000 00000700 0000D4F0 F0F0F1F0 F2F54040 4040D7C6
     0020  E5E3F0F0 F340AB62 537477FC B803C4E2 D5E3C5D7 F340C2C1 E3C3C840 4040C2C1
     0040  E3C3C840 40400000 00000000 0000001A 0001E4E2 C9C2D4E2 E840E2E8 C5C3F1C4
     0060  4040AB62 5362554A 0001

000000010B12  MEMBER(M02     )  URID(000000010A82)  LRSN(AB6253747807)
     TYPE( UNDO )  SUBTYPE(SAVEPOINT)
     *LRH* 002F0090 22000014 0E800000 00010A82 00000001 0A820126 00000001 0A82AB62
           53747807 0002
     0000  00E7D9E4 C9000000 02

000000010B42  MEMBER(M02     )  URID(000000010A82)  LRSN(AB625374780E)  DBID(0113)  OBID(0008)  PAGE(00000003)
     TYPE( UNDO  REDO )  SUBTYPE(TYPE 2 INDEX UPDATE)  CLR(NO)  PROCNAME(DSNKDLE )
     *LRH* 0053002F 06000019 0E800000 00010A82 00000001 0B120126 00000001 0B12AB62
           5374780E 0002
     *LG** 84011300 08000003 63000000 00000000 0000
     0000  001B3000 00B40001 00000201 000A0000 02C5C5F0 F6C1C1D4 F3F1C1
```

Figure 41. Sample DSN1LOGP detail report (Part 1 of 5)

```
000000010B94  MEMBER(M02     )  URID(000000010A82)  LRSN(AB6253747CEF)  DBID(0113)  OBID(0008)  PAGE(00000003)
     TYPE( UNDO  REDO )  SUBTYPE(TYPE 2 INDEX UPDATE)  CLR(NO)  PROCNAME(DSNKINSL)
     *LRH* 00530053 06000019 0E800000 00010A82 00000001 0B420126 00000001 0B42AB62
           53747CEF 0002
     *LG** 04011300 08000003 64000000 00000000 0000
     0000  001B1000 00B30001 00000201 000A2000 00C5C5F0 F6C1C1D7 D7D3F4

......
0000000110A0  MEMBER(M02     )  URID(0000000110A0)  LRSN(AB625379B94A)
     TYPE(UR CONTROL)  SUBTYPE(BEGIN UR)
     *LRH* 009000A0 00200001 03800000 000110A0 00000000 00000126 00000000 0000AB62
           5379B94A 0002
     0000  00020000 00004000 00000000 00000700 0000F0F2 F14BD6D7 D5D3C7D9 F0F0E2E8
     0020  E2D6D7D9 4040AB62 5379B945 07044040 40404040 40400000 00000000 0000E5F4
     0040  F2C44040 40400000 00000000 0000001A 0001E4E2 C9C2D4E2 E840E2E8 C5C3F1C4
     0060  4040AB62 53782F9A 0001

000000011130  MEMBER(M02     )  URID(0000000110A0)  LRSN(AB625379B955)  DBID(0001)  OBID(00CF)  PAGE(00000009)
     TYPE( UNDO  REDO )  SUBTYPE(INSERT IN A DATA PAGE)  CLR(NO)  PROCNAME(DSNISMRT)
     *LRH* 00740090 06000001 0E800000 000110A0 00000001 10A00126 00000001 10A0AB62
           5379B955 0002
     *LG** 90000100 CF000009 3100AB62 51C6F5F1 0000
     0000  003C5036 00D10000 00003400 D1360113 0005F0F7 F1F8F9F5 F0F1F4F5 F2F8F4F8
     0020  00000001 0EB70000 00000000 8000AB62 53782F89 00000000 00000002

0000000111A4  MEMBER(M02     )  LRSN(AB625379D424)  DBID(0001)  OBID(0087)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET OPEN)
     *LRH* 00A00074 00020001 0E800000 00000000 00000001 10000126 00000001 1000AB62
           5379D424 0002
     0000  00010087 6C010100 000000CF 0040C4E2 D5C4C2F0 F140C4E2 D5D3D3E7 F0F10001
     0020  00D10000 10009200 00130020 00000000 00000000 00000000 00000000 00000000
     0040  00000000 00000010 00010000 00000000 00000000 00000000 00000000 00000000
     0060  00AA6F9F 3E74EFAA 80DE586F F5000000 0000C4E2 D5C3F4F1 F040

000000011244  MEMBER(M02     )  URID(0000000110A0)  LRSN(AB625379D47D)  DBID(0001)  OBID(0087)  PAGE(00000005)
     TYPE( UNDO  REDO )  SUBTYPE(TYPE 2 INDEX UPDATE)  CLR(NO)  PROCNAME(DSNKINSL)
     *LRH* 005700A0 06000019 0E800000 000110A0 00000001 11300126 00000001 1130AB62
           5379D47D 0002
     *LG** 00000100 87000005 64000000 00000000 0000
     0000  001F1000 007F0001 00000936 000E2000 00011300 05800000 02549DAC 87D076

0000033592A0  MEMBER(M01     )  LRSN(AB62537A111E)  DBID(0001)  OBID(0086)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET STATUS RECORD)
     *LRH* 00660066 00020009 0E800000 00000000 00000335 923A0126 00000335 923AAB62
           537A111E 0001
     0000  00010086 C4E2D5C4 C2F0F140 C4E2D5D3 D3E7F0F2 00010000 92018000 00000334
     0020  EC3AAB62 5260AB0B 00000000 00000000 00000000 00000000 00000000 00000000

000000011298  MEMBER(M02     )  LRSN(AB62537B40E6)  DBID(0001)  OBID(0086)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET OPEN)
     *LRH* 00A00057 00020001 0E800000 00000000 00000001 11A40126 00000001 11A4AB62
           537B40E6 0002
     0000  00010086 6C010100 000000CF 0040C4E2 D5C4C2F0 F140C4E2 D5D3D3E7 F0F20001
     0020  00D10000 10009200 00130020 00000000 00000000 00000000 00000000 00000000
     0040  00000000 00000010 00010000 00000000 00000000 00000000 00000000 00000000
     0060  00AB624E 7933BEAB 62501309 25000000 0000C4E2 D5C3F4F1 F040
```

Figure 41. Sample DSN1LOGP detail report (Part 2 of 5)

DSN1LOGP

```
00000001133B  MEMBER(M02     )  URID(0000000110A0)  LRSN(AB62537B4242)  DBID(0001)  OBID(0086)  PAGE(00000005)
     TYPE( UNDO  REDO )  SUBTYPE(TYPE 2 INDEX UPDATE)  CLR(NO)  PROCNAME(DSNKINSL)
     *LRH* 005300A0 06000019 0E800000 000110A0 00000001 12440126 00000001 1244AB62
           537B4142 0002
     *LG** 00000100 86000005 64000000 00000000 0000
     0000  001B1000 007F0001 00000936 000A2000 00011300 05AB6253 782F89

00000001138E  MEMBER(M02     )  URID(0000000110A0)  LRSN(AB62537B4931)
     TYPE(UR CONTROL)  SUBTYPE(BEGIN COMMIT1)
     *LRH* 005C0053 00200002 03800000 000110A0 00000001 133B0126 00000001 133BAB62
           537B4931 0002
     0000  00020000 00004000 00000000 00000700 0000F0F2 F14BD6D7 D5D3C7D9 F0F04040
     0020  40404040 40400000 00000000 00000000 00000000 0000

0000000113EA  MEMBER(M02     )  URID(0000000110A0)  LRSN(AB62537B4940)
     TYPE(UR CONTROL)  SUBTYPE(PHASE 1 TO 2)
     *LRH* 0034005C 0020000C 03800000 000110A0 00000001 138E0126 00000001 138EAB62
           537B4940 0002
     0000  00020000 00004000 00000000 0000

0000033685DE  MEMBER(M01     )  LRSN(AB6254D9A231)  DBID(0001)  OBID(001F)
     TYPE( CHECKPOINT)  SUBTYPE(DBE TABLE WITH EXCEPTION DATA)
     *LRH* 0061003E 2100001D 0E800000 00000000 00000336 85A00126 00000336 85A0AB62
           54D9A231 0001
     0000  00000000 C4E2D5C4 C2F0F140 C4C2C4F0 F1404040 0001001F 00000000 00000000
     0020  00000000 00000000 00000000 00000000 00000000 00000000 000000

00000336863F  MEMBER(M01     )  LRSN(AB6254D9A237)  DBID(0001)  OBID(001F)
     TYPE( CHECKPOINT)  SUBTYPE(DBE TABLE WITH PIECE DATA)
     *LRH* 01F60061 2100001E 0E800000 00000000 00000336 85DE0126 00000336 85DEAB62
           54D9A237 0001
     0000  00000100 1FC4E2D5 C4C2F0F1 40C4C2C4 F0F14040 40000000 0020FFFF FFFFFFFF
     0020  00000000 00000000 0000006C 00000090 FFFFFFFF 00000000 00000000 00FFFFFF
     0040  FF000000 00000000 0000FFFF FFFF0000 00000000 000000FF FFFFFF00 00000000
     0060  00000000 FFFFFFFF 00000000 00000000 00FFFFFF FF000000 00000000 0000FFFF
     0080  FFFF0000 00000000 000000FF FFFFFF00 00000000 00000000 FFFFFFFF 00000000
     00A0  00000000 00FFFFFF FF000000 00000000 0000FFFF FFFF0000 00000000 000000FF
     00C0  FFFFFF00 00000000 00000000 FFFFFFFF 00000000 00000000 00FFFFFF FF000000
     00E0  00000000 0000FFFF FFFF0000 00000000 000000FF FFFFFF00 00000000 00000000
     0100  FFFFFFFF 00000000 00000000 00FFFFFF FF000000 00000000 0000FFFF FFFF0000
     0120  00000000 000000FF FFFFFF00 00000000 00000000 FFFFFFFF 00000000 00000000
     0140  00FFFFFF FF000000 00000000 0000FFFF FFFF0000 00000000 000000FF FFFFFF00
     0160  00000000 00000000 FFFFFFFF 00000000 00000000 00FFFFFF FF000000 00000000
     0180  0000FFFF FFFF0000 00000000 000000FF FFFFFF00 00000000 00000000 FFFFFFFF
     01A0  00000000 00000000 00FFFFFF FF000000 00000000 0000FFFF FFFF0000 00000000
     01C0  000000FF FFFFFF00 00000000 00000000

000000003956D  MEMBER(M02     )  LRSN(AB6254EE48E9)  DBID(0006)  OBID(0009)
     TYPE( CHECKPOINT)  SUBTYPE(DBE TABLE WITH EXCEPTION DATA)
     *LRH* 0061003E 2100001D 0E800000 00000000 00000003 952F0126 00000003 952FAB62
           54EE48E9 0002
     0000  00000000 C4E2D5C4 C2F0F640 E2E8E2C4 C2C1E2C5 00060009 00000000 00000000
     0020  00000000 00000000 00000000 00000000 00000000 00000000 000000
```

Figure 41. Sample DSN1LOGP detail report (Part 3 of 5)

```
0000000395CE  MEMBER(M02     )  LRSN(AB6254EE48F2)  DBID(0006)  OBID(0009)
     TYPE( CHECKPOINT)  SUBTYPE(DBE TABLE WITH PIECE DATA)
     *LRH* 01F60061 2100001E 0E800000 00000000 00000003 956D0126 00000003 956DAB62
           54EE48F2 0002
     0000  00000600 09C4E2D5 C4C2F0F6 40E2E8E2 C4C2C1E2 C5000000 0020FFFF FFFFFFFF
     0020  00000000 00000000 000001BC 000001D4 FFFFFFFF 00000000 00000000 00FFFFFF
     0040  FF000000 00000000 0000FFFF FFFF0000 00000000 000000FF FFFFFF00 00000000
     0060  00000000 FFFFFFFF 00000000 00000000 00FFFFFF FF000000 00000000 0000FFFF
     0080  FFFF0000 00000000 000000FF FFFFFF00 00000000 00000000 FFFFFFFF 00000000
     00A0  00000000 00FFFFFF FF000000 00000000 0000FFFF FFFF0000 00000000 000000FF
     00C0  FFFFFF00 00000000 00000000 FFFFFFFF 00000000 00000000 00FFFFFF FF000000
     00E0  00000000 0000FFFF FFFF0000 00000000 000000FF FFFFFF00 00000000 00000000
     0100  FFFFFFFF 00000000 00000000 00FFFFFF FF000000 00000000 0000FFFF FFFF0000
     0120  00000000 000000FF FFFFFF00 00000000 00000000 FFFFFFFF 00000000 00000000
     0140  00FFFFFF FF000000 00000000 0000FFFF FFFF0000 00000000 000000FF FFFFFF00
     0160  00000000 00000000 FFFFFFFF 00000000 00000000 00FFFFFF FF000000 00000000
     0180  0000FFFF FFFF0000 00000000 000000FF FFFFFF00 00000000 00000000 FFFFFFFF
     01A0  00000000 00000000 00FFFFFF FF000000 00000000 0000FFFF FFFF0000 00000000
     01C0  000000FF FFFFFF00 00000000 00000000
  .....

000000057B32  MEMBER(M02     )  LRSN(AB62564FD672)  DBID(0001)  OBID(00CF)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET OPEN)
     *LRH* 00A0006C 00020001 0E800000 00000000 00000005 7AC60126 00000005 7AC6AB62
           564FD672 0002
     0000  000100CF 6C010100 00000000 0040C4E2 D5C4C2F0 F140E2E8 E2D3C7D9 D5E70001
     0020  00000000 10008000 00130020 00000000 00000000 00000000 2A0C0000 00000000
     0040  00000000 00000010 00010000 00000000 00000000 00000000 00000000 00000000
     0060  00AB6251 91E1F9AB 62560CB6 B8000000 0000C4E2 D5C3F4F1 F040

000000057BD2  MEMBER(M02     )  URID(000000057BD2)  LRSN(AB62564FD6D3)
     TYPE(UR CONTROL)  SUBTYPE(BEGIN UR)
     *LRH* 007400A0 00200001 03800000 00057BD2 00000000 00000126 00000000 0000AB62
           564FD6D3 0002
     0000  00070000 00004000 00000000 00000700 0000F0F2 F14BC3D3 E2D3C7D9 F0F0E2E8
     0020  E2D6D7D9 4040AB62 564FD6D0 F7034040 40404040 40400000 00000000 0000E5F4
     0040  F2C44040 40400000 00000000 0000

000000057C46  MEMBER(M02     )  URID(000000057BD2)  LRSN(AB62564FD6DD)  DBID(0001)  OBID(00CF)  PAGE(00000009)
     TYPE( UNDO  REDO )  SUBTYPE(UPDATE NOT IN-PLACE , DATA PART ONLY IN A DATA PAGE)  CLR(NO)  PROCNAME(DSNIREPR)
     *LRH* 006C0074 06000001 0E800000 00057BD2 00000005 7BD20126 00000005 7BD2AB62
           564FD6DD 0002
     *LG** 90000100 CF000009 2D00AB62 54DA2429 0000
     0000  00346137 00D18200 00210013 0013057A 9C8001AB 625391D6 C4AB6256 4FB04700
     0020  02000000 8001AB62 5391D6C4 00000000 00000002

000000057CB2  MEMBER(M02     )  URID(000000057BD2)  LRSN(AB62564FD6FA)
     TYPE(UR CONTROL)  SUBTYPE(BEGIN COMMIT1)
     *LRH* 005C006C 00200002 03800000 00057BD2 00000005 7C460126 00000005 7C46AB62
           564FD6FA 0002
     0000  00070000 00004000 00000000 00000700 0000F0F2 F14BC3D3 E2D3C7D9 F0F04040
     0020  40404040 40400000 00000000 00000000 00000000 0000

000000057D0E  MEMBER(M02     )  URID(000000057BD2)  LRSN(AB62564FD709)
     TYPE(UR CONTROL)  SUBTYPE(PHASE 1 TO 2)
     *LRH* 0034005C 0020000C 03800000 00057BD2 00000005 7CB20126 00000005 7CB2AB62
           564FD709 0002
     0000  00070000 00004000 00000000 0000

000000057D42  MEMBER(M02     )  LRSN(AB62564FE492)  DBID(0001)  OBID(00CF)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET WRITE)
     *LRH* 006C0034 00020007 0E800000 00000000 00000005 7B320126 00000005 7B32AB62
           564FE492 0002
     0000  0000F5E2 C3D40001 00CFC4E2 D5C4C2F0 F140E2E8 E2D3C7D9 D5E70000 00000000
     0020  0000AB62 564FD672 AB62564F D6720000 00000000 01017EFD EAB80000 09000009
     0040  AB62564F D6DD
```

Figure 41. Sample DSN1LOGP detail report (Part 4 of 5)

```
000000057DAE  MEMBER(M02     )  URID(000000057BD2)  LRSN(AB62564FE4C9)
     TYPE(UR CONTROL)  SUBTYPE(END COMMIT2)
     *LRH* 0034006C 00200010 03800000 00057BD2 00000005 7D0E0126 00000005 7D0EAB62
           564FE4C9 0002
     0000  00070000 00004000 00000000 0000

000000057DE2  MEMBER(M02     )  LRSN(AB62564FE550)  DBID(0115)  OBID(0002)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET CLOSE)
     *LRH* 002A0034 00020003 0E800000 00000000 00000005 7D420126 00000005 7D42AB62
           564FE550 0002
     0000  01150002

000000057E0C  MEMBER(M02     )  LRSN(AB62564FEAB1)  DBID(0113)  OBID(000C)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET WRITE)
     *LRH* 006C002A 00020007 0E800000 00000000 00000005 7DE20126 00000005 7DE2AB62
           564FEAB1 0002
     0000  0000F5D7 C3D60113 000CC4C2 C6E6F0F0 F1F1C9E4 C6E6F0F0 F0F30000 00000000
     0020  00000000 0004D98E 00000004 D98E0000 00000000 11010000 00000000 03000003
     0040  AB62553F 9821

000000057E78  MEMBER(M02     )  LRSN(AB62564FF072)  DBID(0113)  OBID(000C)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET CLOSE)
     *LRH* 002A006C 00020003 0E800000 00000000 00000005 7E0C0126 00000005 7E0CAB62
           564FF072 0002
     0000  0113000C

000000057EA2  MEMBER(M02     )  LRSN(AB62564FF606)  DBID(0113)  OBID(000A)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET WRITE)
     *LRH* 009C002A 00020007 0E800000 00000000 00000005 7E780126 00000005 7E78AB62
           564FF606 0002
     0000  0000F5D7 C3D60113 000AC4C2 C6E6F0F0 F1F1C9E4 C6E6F0F0 F0F20000 00000000
     0020  00000000 0004D98E 00000004 D98E0000 00000000 11040000 00000000 03000003
     0040  AB62553F 98780000 00000000 04000004 AB62553F 930C0000 00000000 05000005
     0060  AB62553F 95C30000 00000000 06000006 AB62553F 9855

000000057F3E  MEMBER(M02     )  LRSN(AB62564FFFBA)  DBID(0113)  OBID(000A)
     TYPE(PAGE SET CONTROL)  SUBTYPE(PAGE SET CLOSE)
     *LRH* 002A009C 00020003 0E800000 00000000 00000005 7EA20126 00000005 7EA2AB62
           564FFFBA 0002
     0000  0113000A


DSN1213I DSN1LGRD LAST LOG  RBA ENCOUNTERED 00000337A000


DSN1214I NUMBER OF LOG RECORDS READ 0000000000004661
```

Figure 41. Sample DSN1LOGP detail report (Part 5 of 5)

Interpreting data propagation information in the summary report

The sample output on page Figure 42 on page 607 shows information from the DSN1LOGP summary report about log records of changes to DB2 tables that were defined with DATA CAPTURE CHANGES.

The fields show the following:

- START RBA and END RBA show the first and last RBAs captured for the unit of recovery that was not retrieved. The range that the start and end RBA cover can include one or all of the SQL statements within the scope of the unit of recovery.
- TABLE LIST OVERFLOW tells whether more than 10 distinct data capture table IDs were updated by this unit of recovery. This example shows no overflow occurred.
- LR WRITTEN shows the number of written log records that represented changes to tables defined for data capture and were available to the DB2CDCEX routine. Recursive SQL changes from DB2CDCEX and changes from other attachments

not associated with DB2CDCEX are not included. If you receive a value of 2147483647, an overflow occurred and the count is not valid.

- LR RETRIEVED is the number of captured RBAs that were retrieved by DB2CDCEX. If you receive a value of 2147483647, an overflow occurred and the count is not valid.
- LR NOT RETRIEVED is the difference between the number of written log records (LR WRITTEN) and the number of retrieved log records (LR RETRIEVED). This example shows that four log records were written, and none were retrieved.

```
DATA PROPAGATION INFORMATION:
     START RBA=000004A107F4       END RBA=000004A10A5C           TABLE LIST OVERFLOW=NO
     LR WRITTEN=0000000000000004  LR RETRIEVED=0000000000000000  LR NOT RETRIEVED=0000000000000004
     DATABASE=0112=DBCS1701       PAGESET=0002=TSCS1701          TABLE OBID=0005
```

Figure 42. Sample data propagation information from the summary report

Interpreting error codes

When an error occurs, DSN1LOGP formats a reason code from the DB2 stand-alone log service in the SYSPRINT output. For information about the stand-alone log service and the reason codes it issues, see Appendix C (Volume 2) of *DB2 Administration Guide*.

DSN1LOGP can abend with a user abend code of X'099'. DSN1LOGP finds the corresponding abend reason code in register 15 (at the time of error).

Chapter 38. DSN1PRNT

With the DSN1PRNT stand-alone utility, you can print:

- DB2 VSAM data sets that contain table spaces or index spaces (including dictionary pages for compressed data)
- Image copy data sets
- Sequential data sets that contain DB2 table spaces or index spaces

Using DSN1PRNT, you can print hexadecimal dumps of DB2 data sets and databases. If you specify the FORMAT option, DSN1PRNT formats the data and indexes for any page that does not contain an error that would prevent formatting. If DSN1PRNT detects such an error, it prints an error message just before the page and dumps the page without formatting. Formatting resumes with the next page.

Compressed records are printed in compressed format.

DSN1PRNT is especially useful when you want to identify the contents of a table space or index. You can run DSN1PRNT on image copy data sets as well as table spaces and indexes. DSN1PRNT accepts an index image copy as input when you specify the FULLCOPY option.

DSN1PRNT is compatible with LOB table spaces, when you specify the LOB keyword, and omit the INLCOPY keyword.

Syntax and options of the control statement

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

© Copyright IBM Corp. 1983, 2001

DSN1PRNT syntax diagram

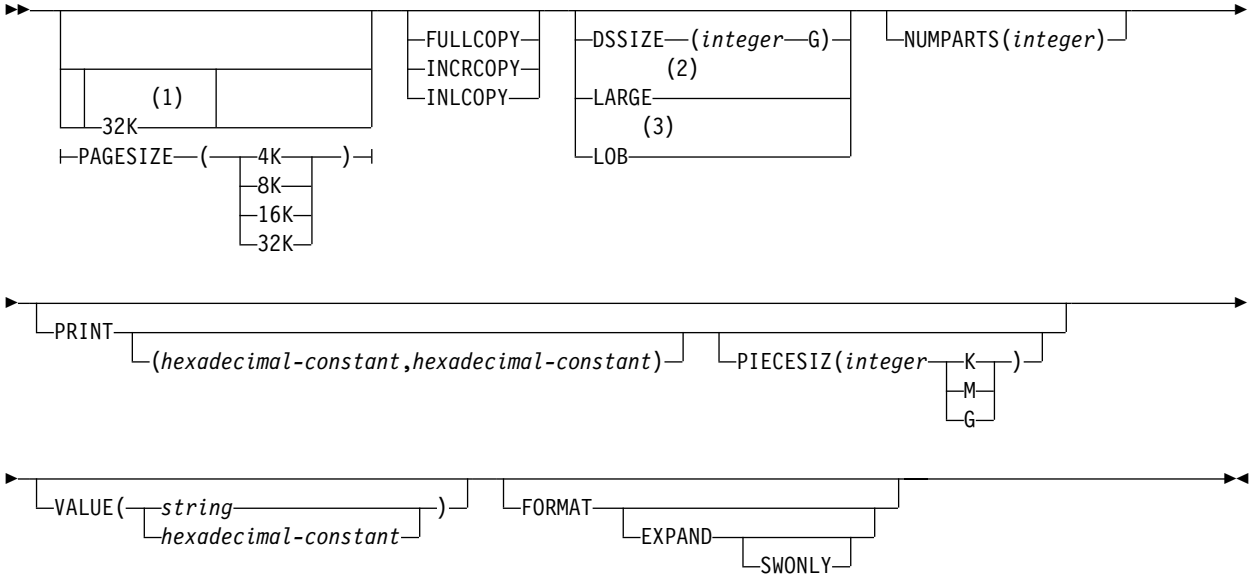

Notes:

1 The preferred option is PAGESIZE(32K).

2 The preferred method of specifying LARGE is DSSIZE(4G).

3 Do not specify INLCOPY with the LOB parameter.

Important: If you specify more than one parameter:
- Separate them by commas (no blanks).
- Specify them in any order.

Option descriptions

Specify one or more of the parameters listed below on the EXEC card to run DSN1PRNT.

32K Specifies that the SYSUT1 data set has a 32-KB page size. If the SYSUT1 data set has a 32-KB page size and you do not specify this option, DSN1PRNT may produce unpredictable results, because the default page size is 4 KB.

The preferred option is **PAGESIZE(32K)**.

PAGESIZE Specifies the page size of the input data set that is defined by SYSUT1. Available page size values are 4K, 8K, 16K, or 32K. If you specify an incorrect page size, DSN1PRNT may produce unpredictable results.

If you omit PAGESIZE, DSN1PRNT tries to determine the page size from the input data set. DB2 issues an error message if DSN1PRNT cannot determine the input page size. This might happen if the header page is not in the input data set, or if the page size field in the header page contains an invalid page size.

DSSIZE(*integer* **G)**

Specifies the data set size, in gigabytes, for the input data set. If you omit DSSIZE, DB2 assumes that the input data set size is 2 GB unless the input data set is a LOB, in which case DB2 assumes a 4 GB input data set size.

integer must match the DSSIZE value specified when the table space was defined.

If you omit DSSIZE and the data set is not one of the default sizes, the results from DSN1PRNT are unpredictable.

LARGE Specifies that the input data set is a table space that was defined with the LARGE option, or an index on such a table space. If you specify LARGE, then DB2 assumes that the data set has a 4-GB boundary.

The preferred method of specifying a table space that was defined with the LARGE option is **DSSIZE(4G)**.

If you omit the LARGE or DSSIZE(4G) option when it is needed, or if you specify LARGE for a table space that was not defined with the LARGE option, the results from DSN1PRNT are unpredictable.

LOB Specifies that the SYSUT1 data set is a LOB table space. You cannot specify the INLCOPY option with the LOB parameter.

DB2 attempts to determine if the input data set is a LOB data set. If you specify the LOB option but the data set is not a LOB data set, or if you omit the LOB option but the data set is a LOB data set, DB2 issues an error message and DSN1PRNT terminates.

NUMPARTS(*integer***)**

Specifies the number of partitions associated with the input data set. NUMPARTS is **required** if the input data set is partitioned. Valid specifications range from 1 to 254. DSN1PRNT uses this value to help locate the first page in a range to be printed. If you omit NUMPARTS or specify it as 0, DSN1PRNT will assume that your input file is not partitioned. If you specify a number greater than 64, DSN1PRNT assumes that the data set is for a partitioned table space that was defined with the LARGE option, even if the LARGE keyword is not specified for DSN1PRNT.

DSN1PRNT cannot always validate the NUMPARTS parameter. If you specify it incorrectly, DSN1PRNT might print the wrong data sets or return an error message indicating that an unexpected page number was encountered.

PRINT(*hexadecimal-constant,hexadecimal-constant***)**

Causes the SYSUT1 data set to be printed in hexadecimal format on the SYSPRINT data set. You can specify the PRINT parameter with or without page range specifications. If you do not specify a range, all pages of the SYSUT1 are printed. If you want to limit the range of pages printed, you can do so by indicating the beginning and ending page numbers with the PRINT parameter or, if you want to print a single page, by indicating only the beginning page. In either case, your range specifications must be from one to eight hexadecimal characters in length.

The following example shows how to code the PRINT parameter if you want to begin printing at page X'2F0' and to stop at page X'35C':

```
PRINT(2F0,35C)
```

To print only the header page for a nonpartitioned table space, specify `PRINT(0)`. For guidance on specifying page numbers for partitioned table spaces, see “Using VERIFY, REPLACE, and DELETE operations” on page 412.

PIECESIZ(*integer***)**

Specifies the maximum piece size (data set size) for non-partitioned indexes. The value you specify must match the value specified when the nonpartitioning index was created or altered.

The defaults for PIECESIZ are 2G (2 GB) for indexes backed by non-large table spaces and 4G (4 GB) for indexes backed by table spaces that were defined with the LARGE option.. This option is required if a print range is specified and the piece size is not one of the default values. If PIECESIZ is omitted and the index is backed by a table space that was defined with the LARGE option, the LARGE keyword is required for DSN1PRNT.

The subsequent keyword K, M, or G, indicates the units of the value specified in *integer*.

K Indicates that the *integer* value is to be multiplied by 1 KB to specify the maximum piece size in bytes. *integer* must be either 256 or 512.

M Indicates that the *integer* value is to be multiplied by 1 MB to specify the maximum piece size in bytes. *integer* must be a power of two, between 1 and 512.

G Indicates that the *integer* value is to be multiplied by 1 GB to specify the maximum piece size in bytes. *integer* must be 1, 2, or 4.

Valid values for piece size are:

- 1 MB or 1 GB
- 2 MB or 2 GB
- 4 MB or 4 GB
- 8 MB
- 16 MB
- 32 MB
- 64 MB
- 128 MB
- 256 KB or 256 MB
- 512 KB or 512 MB

VALUE Causes each page of the input data set SYSUT1 to be scanned for the character string you specify in parentheses following the VALUE parameter. Each page that contains that character string is then printed in SYSPRINT. You can specify the VALUE parameter in conjunction with any of the other DSN1PRNT parameters.

(*string***)**

Can consist of from 1 to 20 alphanumeric characters.

(*hexadecimal-constant***)**
Can consist of 2 to 40 hexadecimal characters. You must specify two single quotation mark characters before and after the hexadecimal character string.

If, for example, you want to search your input file for the string '12345', your JCL might look like the following example:

```
//STEP1 EXEC PGM=DSN1PRNT,PARM='VALUE(12345)'
```

On the other hand, you might want to search for the equivalent hexadecimal character string, in which case your JCL might look like this:

```
//STEP1 EXEC PGM=DSN1PRNT,PARM='VALUE(''F1F2F3F4F5'')'
```

FORMAT Causes the printed output to be formatted. Page control fields are identified and individual records are printed. Empty fields are not displayed.

EXPAND
Specifies that the data is compressed and causes DSN1PRNT to expand it before formatting. This option is intended to be used only under the direction of your IBM Support Center.

SWONLY
Causes DSN1PRNT to use software to expand the compressed data, even when the compression hardware is available. This option is intended to be used only under the direction of your IBM Support Center.

| DSN1PRNT cannot format a leaf or non-leaf page for an index page set that contains keys with altered columns. DSN1PRNT generates message:

```
* KEY WITH ALTERED COLUMN HAS BEEN DETECTED - UNABLE TO FORMAT PAGE *
```

| DSN1PRNT generates unformatted output for the page.

FULLCOPY Specifies that a DB2 full image copy (not a DFSMS concurrent copy) of your data is used as input. If this data is partitioned, you also need to specify the NUMPARTS parameter to identify the number and length of the partitions. If you specify FULLCOPY without including a NUMPARTS specification, DSN1PRNT assumes that the input file is not partitioned.

The FULLCOPY parameter must be specified when using an image copy as input to DSN1PRNT. Omitting the parameter can cause error messages or unpredictable results.

INCRCOPY Specifies that an incremental image copy of the data used as input. If the data is partitioned, also specify NUMPARTS to identify the number and length of the partitions. If you specify INCRCOPY without NUMPARTS, DSN1PRNT assumes that the input file is not partitioned.

The INCRCOPY parameter must be specified when using an incremental image copy as input to DSN1PRNT. Omitting the parameter can cause error messages or unpredictable results.

INLCOPY Specifies that the input data is an inline copy data set.

When you use DSN1PRNT to print a page or a page range from an inline copy produced by LOAD or REORG, DSN1PRNT prints all instances of the pages. The last instance of the page or pages printed is the last one created by the utility.

Before running DSN1PRNT

This section contains information you need to know before you run DSN1PRNT.

Environment

Run DSN1PRNT as an MVS job.

You can run DSN1PRNT even when the DB2 subsystem is not operational. If you choose to use DSN1PRNT when the DB2 subsystem is operational, ensure that the DB2 data sets that are to be printed are not currently allocated to DB2.

To make sure that a data set is not currently allocated to DB2, issue the DB2 STOP DATABASE command, specifying the table spaces and indexes you want to print.

Authorization required

None is required. However, if any of the data sets is RACF-protected, the authorization ID of the job must have RACF authority.

Control statement

See “Syntax and options of the control statement” on page 609 for DSN1PRNT syntax and option descriptions.

Required data sets: DSN1PRNT uses the DD cards described below:

SYSPRINT Defines the data set that contains output messages from DSN1PRNT and all hexadecimal dump output.

SYSUT1 Defines the input data set. That data set can be a sequential data set or a VSAM data set. DSN1PRNT assumes that block size is a multiple of 4096 bytes (as is standard for DB2 data sets).

Disposition for this data set must be specified as OLD (DISP=OLD) to ensure that it is not in use by DB2. Disposition for this data set must be specified as SHR (DISP=SHR) only in circumstances where the DB2 STOP DATABASE command does not work.

The requested operation takes place only for the data set specified. If the input data set belongs to a linear table space or index space that is larger than 2 gigabytes, or if it is a partitioned table space or index space, you must ensure the correct data set is specified. For example, to print a page range in the second partition of a four partition table space, specify NUMPARTS(4) and the data set name of the data set in the group of VSAM data sets comprising the table space. (In other words, `DSN=...A002`.)

If running the online REORG utility with the FASTSWITCH option, verify the data set name before running the DSN1PRNT utility. The fifth level qualifier in the data set name alternates between I0001

and J0001 when using FASTSWITCH. You must specify the correct fifth level qualifier in the data set name to successfully execute the DSN1PRNT utility.

Recommendations

This section contains recommendations for running the DSN1PRNT utility.

Printing with DSN1PRNT instead of DSN1COPY

If you want to print information about a data set, use the DSN1PRNT utility rather than the DSN1COPY utility. This is because DSN1COPY scans the whole SYSUT1 data set, but DSN1PRNT may be able to stop scanning before the end. Also, the DSN1PRNT utility can write a formatted dump.

Determining page size and DSSIZE

Before using DSN1PRNT, determine the page size and data set size (DSSIZE) for the page set. Use the following query on the DB2 catalog to get the information you need:

```
SELECT I.CREATOR,
      I.NAME,
      S.PGSIZE,
      CASE S.DSSIZE
        WHEN 0 THEN CASE S.TYPE
                      WHEN ' ' THEN 2097152
                      WHEN 'I' THEN 2097152
                      WHEN 'L' THEN 4194304
                      WHEN 'K' THEN 4194304
                      ELSE NULL
                    END
        ELSE S.DSSIZE
      END
 FROM SYSIBM.SYSINDEXES I,
      SYSIBM.SYSTABLES  T,
      SYSIBM.SYSTABLESPACE S
 WHERE I.CREATOR='DSN8610' AND
       I.NAME='XEMP1' AND
       I.TBCREATOR=T.CREATOR AND
       I.TBNAME=T.NAME AND
       T.DBNAME=S.DBNAME AND
       T.TSNAME=S.NAME;
```

See “Data sets used by REORG INDEX” on page 316 for information about determining data set names.

Sample control statements

Example 1: Running DSN1PRNT

The fifth qualifier in the data set name can be either I0001 or J0001. Use I0001 in this example.

```
//jobname JOB acct info
//RUNPRNT EXEC PGM=DSN1PRNT,PARM='PRINT,FORMAT'
//STEPLIB DD DSN=prefix.SDSNLOAD
//SYSPRINT DD SYSOUT=A
//SYSUT1 DD DSN=DSNCAT.DSNDBC.DSNDB01.SYSUTIL.I0001.A001,DISP=SHR
```

Example 2: Printing a nonpartitioning index with a 64 MB piece size

The fifth qualifier in the data set name can be either I0001 or J0001. Use I0001 in this example.

```
//PRINT2 EXEC PGM=DSN1PRNT,
//            PARM=(PRINT(F0000,F000F),FORMAT,PIECESIZ(64M))
//* PRINT 16 PAGES IN THE 61ST PIECE OF AN NPI WITH PIECE SIZE OF 64M
//SYSUDUMP DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DISP=OLD,DSN=DSNCAT.DSNDBD.MMRDB.NPI1.I0001.A061
```

DSN1PRNT output

Interpreting output

One intended use of this utility is to aid in determining and correcting system problems. When diagnosing DB2, you might need to refer to licensed documentation to interpret output from this utility.

Chapter 39. DSN1SDMP

Under the direction of the IBM Support Center, use the IFC Selective Dump (DSN1SDMP) utility to:

- Force dumps when selected DB2 trace events occur.
- Write DB2 trace records to a user-defined MVS data set.

For information about the format of trace records, see Appendix D (Volume 2) of *DB2 Administration Guide*.

Syntax and options of the control statement

For guidance in interpreting syntax diagrams, see "How to read the syntax diagrams" on page xviii.

DSN1SDMP syntax diagram

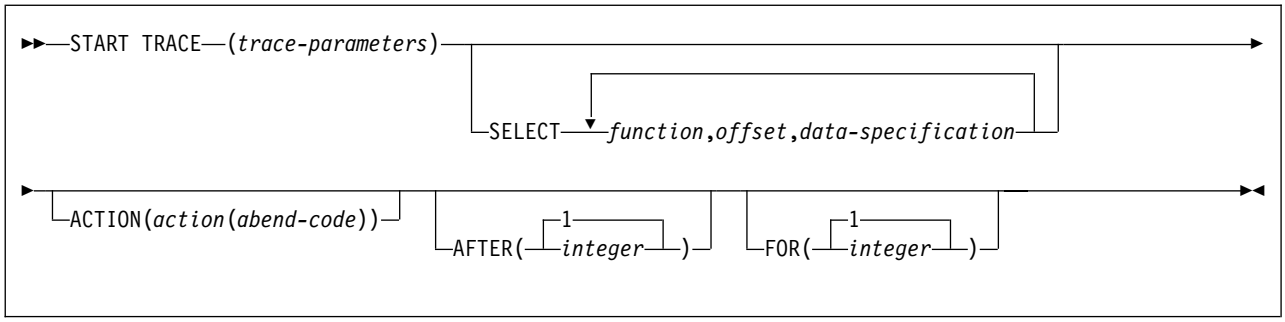

Option descriptions

START TRACE (*trace-parameters***)**

START TRACE is a required keyword and must be the first keyword specified in the SDMPIN input stream. The trace parameters that you use are those described in Chapter 2 of *DB2 Command Reference*, except you do not use the subsystem recognition character.

If the START TRACE command in the SDMPIN input stream is not valid, or the user is not properly authorized, the IFI (instrumentation facility interface) returns an error code and -START TRACE does not take effect. DSN1SDMP writes the error message to the SDMPPRNT data set.

Trace Destination: If DB2 trace data is to be written to the SDMPTRAC data set, the trace destination must be an IFI online performance buffer (OP). OP buffer destinations are specified in the DEST keyword of -START TRACE. There are eight OP buffer destinations, OP1 to OP8. The OPX trace destination assigns the next available OP buffer.

The DB2 output text from the START TRACE command is written to SDMPPRNT.

START TRACE and its associated keywords must be specified first. Specify the remaining selective dump keywords in any order following the START TRACE command.

© Copyright IBM Corp. 1983, 2001

SELECT *function,offset,data-specification*

Specifies selection criteria in addition to those specified on the START TRACE command. SELECT expands the data available to select on in a trace record and allows more specific selection of data in the trace record than using START TRACE alone. A maximum of eight SELECT criteria can be specified.

The selection criteria use the concept of the current-record pointer. The current-record pointer is initialized to zero, meaning the beginning of the trace record. For this instance of the DSN1SDMP trace, the trace record begins with the self-defining section. For information on the fields in the DB2 trace records, see Appendix D (Volume 2) of *DB2 Administration Guide*.

The selection criteria are specified with the following parameters:

function Specifies the type of search to be performed on the trace record. The specified value must be two characters. The possible values are:

DR Direct comparison of data from the specified offset. The offset is always calculated from the current-record pointer.

The current-record pointer locates the point in the trace record where the offset is calculated. The current-record pointer is initialized to zero (the start of the trace record) and is modified by previous P*x* and LN functions (below).

GE Greater than or equal comparison of data from the specified offset. The offset is always calculated from the current-record pointer. The test succeeds if the data from the specified offset is greater than or equal to *data-specification*, which is specified on the SELECT option.

The current-record pointer locates the point in the trace record where the offset is calculated. The current-record pointer is initialized to zero (the start of the trace record) and is modified by previous P*x* and LN functions (below).

LE Less than or equal comparison of data from the specified offset. The offset is always calculated from the current-record pointer. The test succeeds if the data from the specified offset is less than or equal to *data-specification*, which is specified on the SELECT option.

The current-record pointer locates the point in the trace record where the offset is calculated. The current-record pointer is initialized to zero (the start of the trace record) and is modified by previous P*x* and LN functions (below).

P1, P2, or P4

Selects the 1, 2, or 4 byte field located *offset* bytes past the start of the record. Moves the current-record pointer that number of bytes into the record. P1, P2, and P4 always start from the beginning of the record (plus the offset you specify).

This offset is saved as the current-record pointer to be used on subsequent DR, LE, GR, and LN requests.

For example, suppose the user knows the offset to the standard header is 4 bytes long and is located in the first four bytes of the record. P4,00 reads that offset and moves the current-record pointer to the start of the standard header.

LN Advances the current-record pointer by the number of bytes indicated in the 2-byte field located *offset* bytes from the previous current-record pointer.

This offset is saved as the current-record pointer to be used on subsequent DR, LE, GR, and LN requests.

offset A decimal value that specifies the number of bytes into the trace record where the comparison with the *data-specification* field begins. The offset starts from the beginning of the trace record after a P1, P2, or P4, and from the current-record pointer after a GE, LE, LN, or DR.

The format of the DB2 trace record at *data-specification* comparison time is:

Figure 43. Format of the DB2 trace record

- The format of the self-defining section depends on the trace type.
- The format and contents of the data sections depend on the IFCID being recorded. Each record can have one or more data sections. Each data section can have multiple repeating groups.
- The format and content of the trace header section depends on the trace type.

For more information on the format of DB2 trace records, refer to Appendix D (Volume 2) of *DB2 Administration Guide*.

data-specification
Specifies that the data can be hexadecimal (for example, X'9FECBA10') or character (C'FIELD').

ACTION(*action***(***abend-code***))**
Specifies the action to perform when a trace record passes the selection criteria of the START TRACE and SELECT keywords.

Attention: The purpose of the ACTION keyword is to facilitate problem analysis, and it should be used with extreme caution. Not all abends are recoverable, even if the ABENDRET parameter is specified. Some abends may force the DB2 subsystem to terminate, particularly those that occur during end-of-task or end-of-memory processing due to the agent having experienced a previous abend.

action(*abend-code*)
Possible values for *action* are:

ABENDRET ABEND and retry the agent.

ABENDTER ABEND and terminate the agent.

If *action* is not specified, the record is written with no action performed.

An abend reason code can also be specified on this parameter. The codes must be in the range `00E60100-00E60199`. If no abend code is specified, `00E60100` is used.

AFTER(*integer***)**
Specifies that the ACTION is to be performed after the trace point is reached *integer* times.

integer must be between 1 and 32767. The **default** is **AFTER**(**1**).

FOR(*integer***)**
FOR is an optional keyword that specifies the number of times that the ACTION is to take place on reaching the specified trace point. After *integer* times, the trace is stopped and DSN1SDMP terminates.

integer must be between 1 and 32767 and includes the first action. If no SELECT criteria are specified, use an integer greater than 1; the START TRACE command automatically causes the action to take place one time. The **default** is **FOR**(**1**).

Before running DSN1SDMP

This section contains information you need to know before you run DSN1SDMP.

Environment

Run DSN1SDMP as an MVS job and execute it with the DSN TSO command processor. To execute DSN1SDMP, the DB2 subsystem must be running.

The MVS job completes only under the following conditions:

- The TRACE and any additional selection criteria started by DSN1SDMP meet the criteria specified in the FOR parameter.
- The TRACE started by DSN1SDMP is stopped using the STOP TRACE command.
- The job is canceled by the operator.

If you must stop DSN1SDMP, use the STOP TRACE command.

Authorization required

To execute this utility, the privilege set of the process must include one of the following privileges or authorities:

- TRACE system privilege
- SYSOPR authority
- SYSADM authority
- MONITOR1 or MONITOR2 privileges (if you are using user-defined data sets)

The user who executes DSN1SDMP must have EXECUTE authority on the plan specified in the *trace-parameters* of the START TRACE keyword.

Control statement

See “Syntax and options of the control statement” on page 617 for DSN1SDMP syntax and option descriptions.

Required data sets: DSN1SDMP uses the DD cards described below:

SDMPIN Defines the control data set that specifies the input parameters to DSN1SDMP. This DD card is required. The LRECL is 80. Only the first 72 columns are checked by DSN1SDMP.

SDMPPRNT Defines the sequential message data set used for DSN1SDMP messages. If the SDMPPRNT DD statement is omitted, no messages are written. The LRECL is 131.

SYSABEND Defines the data set to contain an ABEND dump in case DSN1SDMP abends. This DD card is optional.

SDMPTRAC Defines the sequential DB2 trace record data set used for trace records returned to DSN1SDMP from DB2. This DD card is required only if trace data is written to an OPX trace destination. If the destination is anything other than an OPX buffer SDMPTRAC is ignored.

Trace records written to SDMPTRAC are of the same format as records written to SMF or GTF, except that, instead of containing the SMF or GTF headers, the SDMPTRAC trace records contain the monitor header (mapped by DSNDQWIW). The DCB parameters are VB, BLKSIZE=8192, LRECL=8188.

SYSTSIN Defines the DSN commands to connect to DB2 and to execute an IFC selective dump:

```
DSN SYSTEM(subsystem name)
RUN PROG(DSN1SDMP) LIB('prefix.SDSNLOAD') PLAN(DSNEDCL)
```

The DB2 subsystem name must be filled in by the user. The DSN RUN command must specify a plan for which the user has execute authority. DSN1SDMP dump does not execute the specified plan; the plan is only used to connect to DB2.

When no plan name is specified on the DSN RUN command, DSN defaults the plan name to the program. When DSN1SDMP is executed without a plan, DSN generates an error if no DSN1SDMP plan exists for which the user has execute authority.

Using DSN1SDMP

This section describes the following tasks associated with running the DSN1SDMP utility:

“Assigning buffers”
“Generating a dump” on page 622
“Stopping or modifying DSN1SDMP traces” on page 622

Assigning buffers

The OPX trace destination assigns the next available OP buffer. You must specify the OPX destination for all traces being recorded to an OP*n* buffer, thereby avoiding the possibility of starting a trace to a buffer that has already been assigned.

If a trace is started to an OP*n* buffer that has already been assigned, DSN1SDMP waits indefinitely until the trace is manually stopped. MONITOR-type traces default to the OPX destination (the next available OP buffer). Other trace types must be explicitly directed to OP destinations via the DEST keyword of the START TRACE command. DSN1SDMP interrogates the IFCAOPN field after the START TRACE COMMAND call to determine if the trace was started to an OP buffer.

Trace records are written to the SDMPTRAC data set when the trace destination is an OP buffer (see page 617). Instrumentation facilities component (IFC) writes trace records to the buffer and posts DSN1SDMP to read the buffer when it fills to half of the buffer size.

You can specify the buffer size on the BUFSIZE keyword of the START TRACE command. The default buffer size is 8KB. All returned records are written to SDMPTRAC.

If the number of generated trace records requires a larger buffer size than was specified, you can lose some trace records. If this happens, you will receive error message DSN2724I.

Generating a dump

All of the following must occur before DSN1SDMP generates a DB2 dump:

- DB2 produces a trace record that satisfies all of the selection criteria.
- An abend action (ABENDRET or ABENDTER) is specified.
- The AFTER and FOR conditions for the trace are satisfied.

If all of these three things occur, an 00E601*xx* abend occurs. *xx* is an integer between 1 and 99 that DB2 obtains from the user-specified value on the ACTION keyword.

Stopping or modifying DSN1SDMP traces

If you must stop DSN1SDMP, use the STOP TRACE command.

If DSN1SDMP does not finish execution, you can stop the utility by issuing the STOP TRACE command; for example:

```
-STOP TRACE=P CLASS(32)
```

DSN1SDMP is designed to execute as a stand-alone batch utility without requiring external intervention from the console operator or other programs. During execution, DSN1SDMP issues an IFI READA request to obtain the data from the OP*n* buffer, and a STOP TRACE command to terminate the original trace started by DSN1SDMP.

A STOP TRACE or MODIFY TRACE command entered from a console against the trace started by DSN1SDMP causes immediate abnormal termination of DSN1SDMP processing. The IFI READA function terminates with an appropriate IFI termination message and reason code. Additional error messages and reason codes associated with the DSN1SDMP STOP TRACE command will vary depending on the specific trace command entered by the console operator.

If the console operator terminates the original trace using the STOP TRACE command, the subsequent STOP TRACE command issued by DSN1SDMP fails.

If the console operator enters a MODIFY TRACE command, the modified trace might also be terminated by the STOP TRACE command issued by DSN1SDMP if MODIFY TRACE processing completes before the DSN1SDMP command is issued.

Sample control statements

Example 1: Skeleton JCL for DSN1SDMP

```
//DSN1J018 JOB 'IFC SD',CLASS=A,
//         MSGLEVEL=(1,1),USER=SYSADM,PASSWORD=SYSADM,REGION=1024K
//*******************************************************************
//*
//*       THIS IS A SKELETON OF THE JCL USED TO RUN DSN1SDMP.
//*       YOU MUST INSERT SDMPIN DD.
//*
//*******************************************************************
//IFCSD    EXEC PGM=IKJEFT01,DYNAMNBR=20,COND=(4,LT)
//STEPLIB  DD DISP=SHR,DSN=prefix.SDSNLOAD
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SDMPPRNT DD SYSOUT=*
//SDMPTRAC DD DISP=(NEW,CATLG,CATLG),DSN=IFCSD.TRACE,
//            UNIT=SYSDA,SPACE=(8192,(100,100)),DCB=(DSORG=PS,
//            LRECL=8188,RECFM=VB,BLKSIZE=8192)
//SDMPIN   DD *
//*******************************************************************
//*
//*       INSERT SDMPIN DD HERE.  IT MUST BEGIN WITH A VALID
//*       START TRACE COMMAND (WITHOUT THE SUBSYSTEM RECOGNITION CHAR)
//*
//*******************************************************************

                  (VALID SDMPIN GOES HERE)

/*
//*******************************************************************
//SYSUDUMP DD SYSOUT=*
//SYSTSIN  DD *
 DSN SYSTEM(DSN)
 RUN PROG(DSN1SDMP) PLAN(DSNEDCL)
 END
//*
```

Example 2: SDMPIN for ABEND and TERMINATE AGENT on -904 SQL CODE

```
//SDMPIN   DD *
* START ONLY IFCID 58, END SQL STATEMENT
 START TRACE=P CLASS(32) IFCID(58) DEST(OPX)
 FOR(1)
 ACTION(ABENDTER(00E60188))
 SELECT
* OFFSET TO FIRST DATA SECTION CONTAINING THE SQLCA.
 P4,08
* SQLCODE -904, RESOURCE UNAVAILABLE
 DR,74,X'FFFFFC78'
/*
```

Example 3: SDMPIN for ABEND and RETRY on RMID 20

```
//*       ABEND AND RETRY AN AGENT WHEN EVENT ID X'0025'
//*       (AGENT ALLOCATION) IS RECORDED BY RMID 20 (SERVICE
//*       CONTROLLER).
//*
//SDMPIN   DD *
*  ENSURE ONLY THE TRACE HEADER IS APPENDED WITH THE STANDARD HEADER
*  VIA THE TDATA KEYWORD ON START TRACE
 START TRACE=P CLASS(3,8) RMID(20) DEST(OPX) TDATA(TRA)
```

```
*  ABEND AND RETRY THE AGENT WITH THE DEFAULT ABEND CODE (00E60100)
 ACTION(ABENDRET)
*  SPECIFY THE SELECT CRITERIA FOR RMID.EID
 SELECT
*  OFFSET TO THE STANDARD HEADER
 P4,00
*  ADD LENGTH OF STANDARD HEADER TO GET TO TRACE HEADER
 LN,00
*  LOOK FOR EID 37 AT OFFSET 4 IN THE TRACE HEADER
 DR,04,X'0025'
/*
```

Example 4: Dump on SQLCODE -811 RMID16 IFCID 58

```
//SDMPIN   DD *
 START TRACE=P CLASS(3) RMID(22) DEST(SMF) TDATA(COR,TRA)
 AFTER(1)
 FOR(1)
 SELECT
* POSITION TO HEADERS (QWHS IS ALWAYS FIRST)
  P4,00
* CHECK QWHS 01, FOR RMID 16, IFCID 58
  DR,02,X'0116003A'
* POSITION TO SECOND SECTION (1ST DATA SECTION)
  P4,08
* COMPARE SQLCODE FOR 811
  DR,74,X'FFFFFCD5'
 ACTION(ABENDRET(00E60188))
/*
```

DSN1SDMP output

Interpreting output

One intended use of this utility is to aid in determining and correcting system problems. When diagnosing DB2, you might need to refer to licensed documentation to interpret output from this utility.

Part 4. Appendixes

© Copyright IBM Corp. 1983, 2001

Appendix A. Limits in DB2 for OS/390 and z/OS

System storage limits might preclude the limits specified here. The limit for items not specified below is system storage.

Table 107. Identifier length limits

Item	Limit
Longest collection ID, correlation name, statement name, or name of an alias, column, cursor, index, table, table check constraint, stored procedure, synonym, user-defined function, view name	18 bytes
Longest authorization name, or name of a database, package, plan, referential constraint, schema, storage group, tablespace name, or trigger	8 bytes
Longest host identifier	64 bytes
Longest server name or location identifier	16 bytes

Table 108. Numeric limits

Item	Limit
Smallest SMALLINT value	-32768
Largest SMALLINT value	32767
Smallest INTEGER value	-2147483648
Largest INTEGER value	2147483647
Smallest REAL value	About -7.2×10^{75}
Largest REAL value	About 7.2×10^{75}
Smallest positive REAL value	About 5.4×10^{-79}
Largest negative REAL value	About -5.4×10^{-79}
Smallest FLOAT value	About -7.2×10^{75}
Largest FLOAT value	About 7.2×10^{75}
Smallest positive FLOAT value	About 5.4×10^{-79}
Largest negative FLOAT value	About -5.4×10^{-79}
Smallest DECIMAL value	$1 - 10^{31}$
Largest DECIMAL value	$10^{31} - 1$
Largest decimal precision	31

Table 109. String length limits

Item	Limit
Maximum length of CHAR	255 bytes
Maximum length of GRAPHIC	127 DBCS characters
Maximum length[6] of VARCHAR	4046 bytes for 4-KB pages 8128 bytes for 8-KB pages 16320 bytes for 16-KB pages 32704 bytes for 32-KB pages

© Copyright IBM Corp. 1983, 2001

Table 109. String length limits (continued)

Item	Limit
Maximum length[6] of VARGRAPHIC	4046 bytes (2023 DBCS characters) for 4-KB pages 8128 bytes (4064 DBCS characters) for 8-KB pages 16320 bytes (8160 DBCS characters for 16-KB pages 32704 bytes (16352 DBCS characters) for 32-KB pages
Maximum length of CLOB	2 147 483 647 bytes (2 gigabytes - 1 byte)
Maximum length of DBCLOB	1 073 741 824 DBCS characters
Maximum length of BLOB	2 147 483 647 bytes (2 gigabytes - 1 byte)
Maximum length of a character constant	255 bytes
Maximum length of a hexadecimal constant	254 digits
Maximum length of a graphic string constant	124 DBCS characters
Maximum length of a concatenated character string	2 147 483 647 bytes (2 gigabytes - 1 byte)
Maximum length of a concatenated graphic string	1 073 741 824 DBCS characters
Maximum length of a concatenated binary string	2 147 483 647 bytes (2 gigabytes - 1 byte)

Table 110. Datetime limits

Item	Limit
Smallest DATE value (shown in ISO format)	0001-01-01
Largest DATE value (shown in ISO format)	9999-12-31
Smallest TIME value (shown in ISO format)	00.00.00
Largest TIME value (shown in ISO format)	24.00.00
Smallest TIMESTAMP value	0001-01-01-00.00.00.000000
Largest TIMESTAMP value	9999-12-31-24.00.00.000000

Table 111. DB2 limits on SQL statements

Item	Limit
Maximum number of columns in a table or view (the value depends on the complexity of the CREATE VIEW statement) or columns returned by a table function.	750 or fewer 749 if the table is a dependent
Maximum number of base tables in a view, SELECT, UPDATE, INSERT, or DELETE	225
Maximum row and record sizes for a table	See the description of CREATE TABLE in Chapter 5 of *DB2 SQL Reference*
Maximum number of volume IDs in a storage group	133
Maximum number of partitions in a partitioned table space or partitioned index	64 for table spaces that are not defined with LARGE or a DSSIZE greater than 2G 254 for table spaces that are defined with LARGE or a DSSIZE greater than 2G

6. The maximum length can be achieved only if the column is the only column in the table. Otherwise, the maximum length depends on the amount of space remaining on a page.

Table 111. DB2 limits on SQL statements (continued)

Item	Limit
Maximum size of a partition (table space or index)	For table spaces that are not defined with LARGE or a DSSIZE greater than 2G: 4 gigabytes, for 1 to 16 partitions 2 gigabytes, for 17 to 32 partitions 1 gigabyte, for 33 to 64 partitions For table spaces that are defined with LARGE: 4 gigabytes, for 1 to 254 partitions For table spaces that are defined with a DSSIZE greater than 2G: 64 gigabytes, for 1 to 254 partitions
Maximum size of a DBRM entry	131072 bytes
Longest index key	255 bytes less the number of key columns that allow nulls.
Maximum number of bytes used in the partitioning of a partitioned index[7]	255
Maximum number of columns in an index key	64
Maximum number of tables in a FROM clause	15
Maximum number of subqueries in a statement	14
Maximum total length of host and indicator variables pointed to in an SQLDA	32767 bytes 2 147 483 647 bytes (2 gigabytes - 1 byte) for a LOB, subject to the limitations imposed by the application environment and host language
Longest host variable used for insert or update	32704 bytes for a non-LOB 2 147 483 647 bytes (2 gigabytes - 1 byte) for a LOB, subject to the limitations imposed by the application environment and host language
Longest SQL statement	32765 bytes
Maximum number of elements in a select list	750
Maximum number of predicates in a WHERE or HAVING clause	750
Maximum total length of columns of a query operation requiring a sort key (SELECT DISTINCT, ORDER BY, GROUP BY, UNION without the ALL keyword, and the DISTINCT column function)	4000 bytes
Maximum length of a table check constraint	3800 bytes
Maximum number of bytes that can be passed in a single parameter of an SQL CALL statement	32765 bytes for a non-LOB 2 147 483 647 bytes (2 gigabytes - 1 byte) for a LOB, subject to the limitations imposed by the application environment and host language
Maximum number of stored procedures, triggers, and user-defined functions that an SQL statement can implicitly or explicitly reference	16 nesting levels
Maximum length of the SQL path	254 bytes

7. The maximum length of the key for a partitioning index is 255 bytes; all 255 bytes can be used to determine the partition.

Limits in DB2 for OS/390 and z/OS

Table 112. DB2 system limits

Item	Limit
Maximum number of concurrent DB2 or application agents	Limited by the EDM pool size, buffer pool size, and the amount of storage used by each DB2 or application agent
Largest table or table space	16 terabytes
Largest log space	2^{48}
Largest active log data set	4 gigabytes -1
Largest archive log data set	4 gigabytes -1
Maximum number of active log copies	2
Maximum number of archive log copies	2
Maximum number of active log data sets (each copy)	31
Maximum number of archive log volumes (each copy)	1000
Maximum number of databases accessible to an application or end user	Limited by system storage and EDM pool size
Largest EDM pool	The installation parameter maximum depends on available space
Maximum number of databases	65279
Maximum number of rows per page	255 for all table spaces except catalog and directory tables spaces, which have a maximum of 127
Maximum simple or segmented data set size	2 gigabytes
Maximum partitioned data set size	See item “maximum size of a partition” in Table 111 on page 628
Maximum LOB data set size	64 gigabytes

Appendix B. Stored procedures shipped with DB2

DB2 provides several stored procedures that you can call in your application programs to perform a number of utility and application programming functions. Those stored procedures are:

- The utilities stored procedure (DSNUTILS)

  This stored procedure lets you invoke utilities from a local or remote client program. See “Invoking utilities as a stored procedure (DSNUTILS)” for information.
- The DB2 UDB Control Center table space and index information stored procedure (DSNACCQC)

  This stored procedure helps you determine when utilities should be run on your databases. This stored procedure is designed primarily for use by the DB2 UDB Control Center but can be invoked from any client program. See “The DB2 UDB Control Center table space and index information stored procedure (DSNACCQC)” on page 642 for information.
- The DB2 UDB Control Center partition information stored procedure (DSNACCAV)

  This stored procedure helps you determine when utilities should be run on your partitioned table spaces. This stored procedure is designed primarily for use by the DB2 UDB Control Center but can be invoked from any client program. See “The DB2 UDB Control Center partition information stored procedure (DSNACCAV)” on page 649 for information.
- The real-time statistics stored procedure (DSNACCOR)

  This stored procedure queries the DB2 real-time statistics tables to help you determine when you should run COPY, REORG, or RUNSTATS, or enlarge your DB2 data sets. See “The DB2 real-time statistics stored procedure” on page 658 for information.
- The WLM environment refresh stored procedure (WLM_REFRESH)

  This stored procedure lets you refresh a WLM environment from a remote workstation. See Appendix I of *DB2 Application Programming and SQL Guide* for information.
- The CICS transaction invocation stored procedure (DSNACICS)

  This stored procedure lets you invoke CICS transactions fom a remote workstation. See Appendix I of *DB2 Application Programming and SQL Guide* for information.

Invoking utilities as a stored procedure (DSNUTILS)

The DSNUTILS stored procedure enables you use the SQL CALL statement to execute DB2 utilities from a DB2 application program. When called, DSNUTILS performs the following actions:

- Dynamically allocates the specified data sets
- Creates the utility input (SYSIN) stream
- Invokes DB2 utilities (Program DSNUTILB)
- Deletes all the rows currently in the created temporary table (SYSIBM.SYSPRINT)
- Captures the utility output stream (SYSPRINT) into a created temporary table (SYSIBM.SYSPRINT)
- Declares a cursor to select from SYSPRINT:

© Copyright IBM Corp. 1983, 2001

```
DECLARE SYSPRINT CURSOR WITH RETURN FOR
  SELECT SEQNO, TEXT FROM SYSPRINT
  ORDER BY SEQNO;
```

- Opens the SYSPRINT cursor and returns.

The calling program then fetches from the returned result set to obtain the captured utility output.

Environment

DSNUTILS ***must*** run in a WLM environment. TCB=1 is also required.

Authorization required

To execute the CALL statement, the owner of the package or plan that contains the CALL statement must have one or more of the following privileges on each package that the stored procedure uses:

- The EXECUTE privilege on the package for DSNUTILS
- Ownership of the package
- PACKADM authority for the package collection
- SYSADM authority

Then, to execute the utility, the privilege set must also include the authorization to run the specified utility.

Control statement

DSNUTILS dynamically allocates the specified data sets. Any utility that requires a sort must include the SORTDEVT keyword in the utility control statement, and optionally, the SORTNUM keyword.

If the DSNUTILS stored procedure invokes a new utility, refer to Table 113 for information about the default data dispositions specified for dynamically allocated data sets.

Table 113. Data dispositions for dynamically allocated data sets

ddname	CHECK DATA	CHECK INDEX or CHECK LOB	COPY	COPY-TOCOPY	LOAD	MERGE-COPY	REBUILD INDEX	REORG INDEX	REORG TABLE-SPACE	UNLOAD
SYSREC	ignored	ignored	ignored	ignored	OLD KEEP KEEP	ignored	ignored	ignored	NEW CATLG CATLG	NEW CATLG CATLG
SYSDISC	ignored	ignored	ignored	ignored	NEW CATLG CATLG	ignored	ignored	ignored	NEW CATLG CATLG	ignored
SYSPUNCH	ignored	ignored	ignored	ignored	ignored	ignored	ignored	ignored	NEW CATLG CATLG	NEW CATLG CATLG
SYSCOPY	ignored	ignored	NEW CATLG CATLG	ignored	NEW CATLG CATLG	NEW CATLG CATLG	ignored	ignored	NEW CATLG CATLG	ignored
SYSCOPY2	ignored	ignored	NEW CATLG CATLG	ignored	NEW CATLG CATLG	NEW CATLG CATLG	ignored	ignored	NEW CATLG CATLG	ignored

Table 113. Data dispositions for dynamically allocated data sets (continued)

ddname	CHECK DATA	CHECK INDEX or CHECK LOB	COPY	COPY-TOCOPY	LOAD	MERGE-COPY	REBUILD INDEX	REORG INDEX	REORG TABLE-SPACE	UNLOAD
SYSRCPY1	ignored	ignored	NEW CATLG CATLG	ignored	NEW CATLG CATLG	NEW CATLG CATLG	ignored	ignored	NEW CATLG CATLG	ignored
SYSRCPY2	ignored	ignored	NEW CATLG CATLG	ignored	NEW CATLG CATLG	NEW CATLG CATLG	ignored	ignored	NEW CATLG CATLG	ignored
SYSUT1	NEW DELETE CATLG	NEW DELETE CATLG	ignored	ignored	NEW DELETE CATLG	ignored	NEW DELETE CATLG	NEW CATLG CATLG	NEW DELETE CATLG	ignored
SORTOUT	NEW DELETE CATLG	ignored	ignored	ignored	NEW DELETE CATLG	ignored	ignored	ignored	NEW DELETE CATLG	ignored
SYSMAP	ignored	ignored	ignored	ignored	NEW CATLG CATLG	ignored	ignored	ignored	ignored	ignored
SYSERR	NEW CATLG CATLG	ignored	ignored	ignored	NEW CATLG CATLG	ignored	ignored	ignored	ignored	ignored
FILTER	ignored	ignored	NEW DELETE CATLG	ignored	ignored	ignored	ignored	ignored	ignored	ignored

If the DSNUTILS stored procedure restarts a current utility, refer to Table 114 for information about the default data dispositions specified for dynamically-allocated data sets.

Table 114. Data dispositions for dynamically allocated data sets on RESTART

ddname	CHECK DATA	CHECK INDEX or CHECK LOB	COPY	COPY-TOCOPY	LOAD	MERGE-COPY	REBUILD INDEX	REORG INDEX	REORG TABLE-SPACE	UNLOAD
SYSREC	ignored	ignored	ignored	ignored	OLD KEEP KEEP	ignored	ignored	ignored	MOD CATLG CATLG	MOD CATLG CATLG
SYSDISC	ignored	ignored	ignored	ignored	MOD CATLG CATLG	ignored	ignored	ignored	MOD CATLG CATLG	ignored
SYSPUNCH	ignored	ignored	ignored	ignored	ignored	ignored	ignored	ignored	MOD CATLG CATLG	MOD CATLG CATLG
SYSCOPY	ignored	ignored	MOD CATLG CATLG	ignored	MOD CATLG CATLG	MOD CATLG CATLG	ignored	ignored	MOD CATLG CATLG	ignored
SYSCOPY2	ignored	ignored	MOD CATLG CATLG	ignored	MOD CATLG CATLG	MOD CATLG CATLG	ignored	ignored	MOD CATLG CATLG	ignored

Table 114. Data dispositions for dynamically allocated data sets on RESTART (continued)

ddname	CHECK DATA	CHECK INDEX or CHECK LOB	COPY	COPY-TOCOPY	LOAD	MERGE-COPY	REBUILD INDEX	REORG INDEX	REORG TABLE-SPACE	UNLOAD
SYSRCPY1	ignored	ignored	MOD CATLG CATLG	ignored	MOD CATLG CATLG	MOD CATLG CATLG	ignored	ignored	MOD CATLG CATLG	ignored
SYSRCPY2	ignored	ignored	MOD CATLG CATLG	ignored	MOD CATLG CATLG	MOD CATLG CATLG	ignored	ignored	MOD CATLG CATLG	ignored
SYSUT1	MOD DELETE CATLG	MOD DELETE CATLG	ignored	ignored	MOD DELETE CATLG	ignored	MOD DELETE CATLG	MOD CATLG CATLG	MOD DELETE CATLG	ignored
SORTOUT	MOD DELETE CATLG	ignored	ignored	ignored	MOD DELETE CATLG	ignored	ignored	ignored	MOD DELETE CATLG	ignored
SYSMAP	ignored	ignored	ignored	ignored	MOD CATLG CATLG	ignored	ignored	ignored	ignored	ignored
SYSERR	MOD CATLG CATLG	ignored	ignored	ignored	MOD CATLG CATLG	ignored	ignored	ignored	ignored	ignored
FILTER	ignored	ignored	MOD DELETE CATLG	ignored	ignored	ignored	ignored	ignored	ignored	ignored

DSNUTILS syntax diagram

For guidance in interpreting syntax diagrams, see “How to read the syntax diagrams” on page xviii.

The following syntax diagram shows the SQL CALL statement for invoking utilities as a stored procedure.

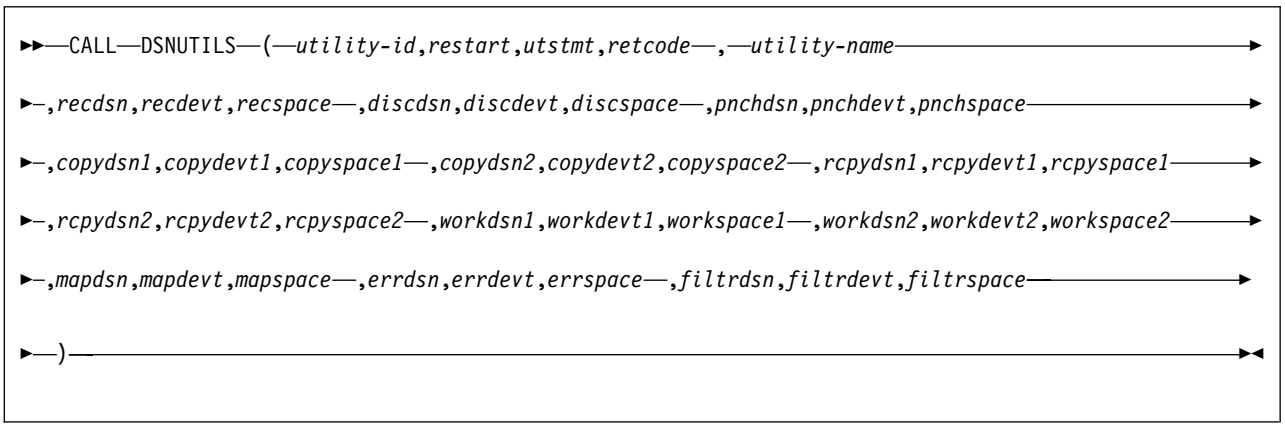

DSNUTILS option descriptions

utility-id
Specifies a unique identifier for this utility within DB2.

This is an input parameter of type VARCHAR(16).

restart Specifies whether this restarts a current utility, and, if so, at what point it is to be restarted.

This is an input parameter of type VARCHAR(8).

NO or null
Indicates the utility is new, not a restart. There must not be any other utility with the same utility identifier (UID).

The **default** is **null**.

CURRENT
Restarts the utility at the last commit point.

PHASE
Restarts the utility at the beginning of the currently stopped phase. Use the DISPLAY UTILITY to determine the currently stopped phase.

PREVIEW
Executes in PREVIEW mode the utility control statements that follow. It parses for syntax errors all utility control statements, but normal utility execution will not take place. If syntax is valid, the utility will expand all LISTDEF lists and TEMPLATE DSNAMEs that appear in SYSIN and print results to the SYSPRINT data set. PREVIEW will evaluate and expand all LISTDEFs into an actual list of table spaces or index spaces. It evaluates TEMPLATE DSNAMEs into actual data set names through variable substitution. It also expands lists from the SYSLISTD DD and TEMPLATE DSNAMEs from the SYSTEMPL DD that is referenced by a utility invocation.

Absence of the PREVIEW keyword turns off preview processing with one exception. That it does not over ride the PREVIEW JCL parameter which, if specified, remains in affect for the entire job step.

This option is identical to the PREVIEW JCL parameter.

utstmt Specifies the utility control statements.

This is an input parameter of type VARCHAR(32704).

retcode
Specifies the utility highest return code.

This is an output parameter of type INTEGER.

utility-name
Specifies the utility you want to invoke. Specify a value of **ANY** if you are using TEMPLATE dynamic allocation. This will thwart redundant dynamic allocation by DSNUTILS.

This is an input parameter of type VARCHAR(20).

Note: Because it allows only a single utility here, there is limited dynamic support of data set allocation. Specify only a single utility requiring data set allocation in the *utstmt* parameter.

Select the utility name from the following list:

ANY*
CHECK DATA

CHECK INDEX
CHECK LOB
COPY
DIAGNOSE
LOAD
MERGECOPY
MODIFY RECOVERY
MODIFY STATISTICS
QUIESCE
REBUILD INDEX
RECOVER
REORG INDEX
REORG LOB
REORG TABLESPACE
REPAIR
REPORT RECOVERY
REPORT TABLESPACESET
RUNSTATS INDEX
RUNSTATS TABLESPACE
STOSPACE
UNLOAD

Note: Use ANY to indicate that TEMPLATE dynamic allocation will be used. This suppresses the dynamic allocation normally performed by DSNUTILS.

recdsn Specifies the cataloged data set name required by LOAD for input, or by REORG TABLESPACE as the unload data set. *recdsn* is required for LOAD. It is also required for REORG TABLESPACE unless you also specified NOSYSREC or SHRLEVEL CHANGE. If you specify *recdsn*, it will be allocated to the **SYSREC** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **INDDN** parameter for LOAD, the value specified for *ddname* **MUST** be **SYSREC**.

If you specified the **UNLDDN** parameter for REORG TABLESPACE, the value specified for *ddname* **MUST** be **SYSREC**.

recdevt
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *recdsn* data set resides.

This is an input parameter of type CHAR(8).

recspace
Specifies the number of cylinders to use as the primary space allocation for the *recdsn* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

discdsn
Specifies the cataloged data set name used by LOAD as a discard data set to hold records not loaded, and by REORG TABLESPACE as a discard data set to hold records not reloaded. If you specify *discdsn*, it will be allocated to the **SYSDISC** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **DISCARDDN** parameter for LOAD or REORG TABLESPACE, the value specified for *ddname* **MUST** be **SYSDISC**.

discdevt
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *discdsn* data set resides.

This is an input parameter of type CHAR(8).

discspace
Specifies the number of cylinders to use as the primary space allocation for the *discdsn* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

pnchdsn
Specifies the cataloged data set name that REORG TABLESPACE UNLOAD EXTERNAL or REORG TABLESPACE DISCARD uses to hold the generated LOAD utility control statements. If you specify a value for *pnchdsn*, it will be allocated to the **SYSPUNCH** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **PUNCHDDN** parameter for REORG TABLESPACE, the value specified for *ddname* **MUST** be **SYSPUNCH**.

pnchdevt
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *pnchdsn* data set resides.

This is an input parameter of type CHAR(8).

pnchspace
Specifies the number of cylinders to use as the primary space allocation for the *pnchdsn* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

copydsn1
Specifies the name of the required target (output) data set, which is needed when you specify the COPY utility or the MERGECOPY utility. It is optional for LOAD and REORG TABLESPACE. If you specify *copydsn1*, it will be allocated to the **SYSCOPY** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **COPYDDN** parameter for COPY, MERGECOPY, LOAD, or REORG TABLESPACE, the value specified for *ddname1* **MUST** be **SYSCOPY**.

copydevt1
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *copydsn1* data set resides.

This is an input parameter of type CHAR(8).

copyspace1
Specifies the number of cylinders to use as the primary space allocation for the *copydsn1* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

copydsn2
Specifies the name of the cataloged data set used as a target (output) data set for the backup copy. It is optional for COPY, MERGECOPY, LOAD, and REORG TABLESPACE. If you specify *copydsn2*, it will be allocated to the **SYSCOPY2** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **COPYDDN** parameter for COPY, MERGECOPY, LOAD, or REORG TABLESPACE, the value specified for *ddname2* **MUST** be **SYSCOPY2**.

copydevt2
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *copydsn2* data set resides.

This is an input parameter of type CHAR(8).

copyspace2
Specifies the number of cylinders to use as the primary space allocation for the *copydsn2* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

rcpydsn1
Specifies the name of the cataloged data set required as a target (output) data set for the remote site primary copy. It is optional for COPY, LOAD, and REORG TABLESPACE. If you specified *rcpydsn1*, it will be allocated to the **SYSRCPY1** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **RECOVERYDDN** parameter for COPY, MERGECOPY, LOAD, or REORG TABLESPACE, the value specified for *ddname1* **MUST** be **SYSRCPY1**.

rcpydevt1
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *rcpydsn1* data set resides.

This is an input parameter of type CHAR(8).

rcpyspace1
Specifies the number of cylinders to use as the primary space allocation for the *rcpydsn1* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

rcpydsn2
Specifies the name of the cataloged data set required as a target (output) data set for the remote site backup copy. It is optional for COPY, LOAD, and REORG TABLESPACE. If you specify *rcpydsn2*, it will be allocated to the **SYSRCPY2** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **RECOVERYDDN** parameter for COPY, MERGECOPY, LOAD, or REORG TABLESPACE, the value specified for *ddname2* **MUST** be **SYSRCPY2**.

rcpydevt2
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *rcpydsn2* data set resides.

This is an input parameter of type CHAR(8).

rcpyspace2
Specifies the number of cylinders to use as the primary space allocation for the *rcpydsn2* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

workdsn1
Specifies the name of the cataloged data set required as a work data set for sort input and output. It is required for CHECK DATA, CHECK INDEX and REORG INDEX. It is also required for LOAD and REORG TABLESPACE unless you also specified the SORTKEYS keyword. It is optional for REBUILD INDEX. If you specify *workdsn1*, it will be allocated to the **SYSUT1** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **WORKDDN** parameter for CHECK DATA, CHECK INDEX, LOAD, REORG INDEX, REORG TABLESPACE, or REBUILD INDEX, the value specified for *ddname* **MUST** be **SYSUT1**.

workdevt1
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *workdsn1* data set resides.

This is an input parameter of type CHAR(8).

workspace1
Specifies the number of cylinders to use as the primary space allocation for the *workdsn1* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

workdsn2
Specifies the name of the cataloged data set required as a work data set for sort input and output. It is required for CHECK DATA. It is also required if you are using REORG INDEX to reorganize non-unique type 1 indexes. It is required for LOAD or REORG TABLESPACE unless you also specified the SORTKEYS keyword. If you specify *workdsn2*, it will be allocated to the **SORTOUT** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **WORKDDN** parameter for CHECK DATA, LOAD, REORG INDEX, or REORG TABLESPACE, the value specified for *ddname* **MUST** be **SORTOUT**.

workdevt2
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *workdsn2* data set resides.

This is an input parameter of type CHAR(8).

workspace2
Specifies the number of cylinders to use as the primary space allocation for the *workdsn2* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

mapdsn
Specifies the name of the cataloged data set required as a work data set for error processing during LOAD with ENFORCE CONSTRAINTS. It is optional for LOAD. If you specify *mapdsn*, it will be allocated to the **SYSMAP** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **MAPDDN** parameter for LOAD, the value specified for *ddname* **MUST** be **SYSMAP**.

mapdevt
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *mapdsn* data set resides.

This is an input parameter of type CHAR(8).

mapspace
Specifies the number of cylinders to use as the primary space allocation for the *mapdsn* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

errdsn Specifies the name of the cataloged data set required as a work data set for error processing. It is required for CHECK DATA, and is optional for LOAD. If you specify *errdsn*, it will be allocated to the **SYSERR** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **ERRDDN** parameter for CHECK DATA or LOAD, the value specified for *ddname* **MUST** be **SYSERR**.

errdevt
Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *errdsn* data set resides.

This is an input parameter of type CHAR(8).

errspace
Specifies the number of cylinders to use as the primary space allocation for the *errdsn* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

filtrdsn Specifies the name of the cataloged data set required as a work data set for error processing. It is optional for COPY CONCURRENT. If you specify *filtrdsn*, it will be allocated to the **FILTER** DDNAME.

This is an input parameter of type VARCHAR(54).

Note: If you specified the **FILTERDDN** parameter for COPY, the value specified for *ddname* **MUST** be **FILTER**.

filtrdevt

Specifies a unit address, a generic device type, or a user-assigned group name for a device on which the *filtrdsn* data set resides.

This is an input parameter of type CHAR(8).

filtrspace

Specifies the number of cylinders to use as the primary space allocation for the *filtrdsn* data set. The secondary space allocation will be 10% of the primary.

This is an input parameter of type SMALLINT.

Modifying the WLM-established address space

Add UTPRINT, DSSPRINT, SYSIN and SYSPRINT to the JCL procedure for starting the WLM-established address space, in which DSNUTILS runs. You must allocate SYSIN and SYSPRINT in the procedure to temporarily store utility input statements and utility output messages. Use JCL similar to the following sample PROC:

```
//*************************************************************
//*    JCL FOR RUNNING THE WLM-ESTABLISHED STORED PROCEDURES
//*    ADDRESS SPACE
//*       RGN     -- THE MVS REGION SIZE FOR THE ADDRESS SPACE.
//*       DB2SSN  -- THE DB2 SUBSYSTEM NAME.
//*       APPLENV -- THE MVS WLM APPLICATION ENVIRONMENT
//*                  SUPPORTED BY THIS JCL PROCEDURE.
//*
//*************************************************************
//DSNWLM   PROC RGN=0K,APPLENV=WLMENV1,DB2SSN=DSN
//IEFPROC EXEC PGM=DSNX9WLM,REGION=&RGN,TIME=NOLIMIT,
//        PARM='&DB2SSN,1,&APPLENV'
//STEPLIB  DD  DISP=SHR,DSN=CEE.V!R!M!.SCEERUN
//         DD  DISP=SHR,DSN=DSN!!0.SDSNLOAD
//UTPRINT  DD  SYSOUT=*
//DSSPRINT DD  SYSOUT=*
//SYSIN    DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
//SYSPRINT DD  UNIT=SYSDA,SPACE=(4000,(20,20),,,ROUND)
```

Sample program for calling DSNUTILS

Three example programs calling DSNUTILS are shipped in SDSNSAMP.

- DSNTEJ6U - A DSNUTILS caller using PL/I. Job DSNTEJ6U compiles, link-edits, binds, and runs sample PL/I program DSN8EPU, which invokes the DSNUTILS stored procedure to execute an utility.
- DSNTEJ6V - A DSNUTILS caller using C++. Job DSNTEJ6V compiles, link-edits, binds, and runs sample C++ program DSN8EE1, which invokes the DSNUTILS stored procedure to execute an utility.
- DSNTEJ80 - A DSNUTILS caller using C and ODBC. You can use this sample to compile, pre-link, link-edit, and execute the sample application DSN8OIVP, which can be used to verify that your DB2 ODBC installation is correct.

DSNUTILS output

DB2 creates the result set according to the DECLARE statement shown on page 631.

Output from a successful execution of the DSNTEJ6U sample job or an equivalent job lists the parameters specified followed by the messages generated by the DB2 DIAGNOSE DISPLAY MEPL utility.

The DB2 UDB Control Center table space and index information stored procedure (DSNACCQC)

The information under this heading is Product-sensitive Programming Interface and Associated Guidance Information, as defined in “Notices” on page 725.

DSNACCQC is a sample stored procedure that gives you information about your table spaces and indexes. You can use DSNACCQC to obtain the following types of information:

- Table spaces and indexes on which RUNSTATS needs to be run
- Table spaces and indexes on which the STOSPACE utility has not been run
- Table spaces and indexes that exceed the primary space allocation
- Table spaces with more than a user-specified percentage of relocated rows
- Table spaces with more than a user-specified percentage of space that is occupied by dropped tables
- Table spaces with table space locking
- Simple table spaces with more than one table
- Indexes with clustering problems
- Indexes with more than a user-specified number of index levels
- Indexes with more than a user-specified LEAFDIST value
- Type 1 indexes
- Nonunique indexes with long RID chains
- Indexes that are not used in static SQL statements

Environment

DSNACCQC runs in a WLM-established stored procedures address space.

Authorization required

To execute the CALL statement, the owner of the package or plan that contains the CALL statement must have one or more of the following privileges on each package that the stored procedure uses:

- The EXECUTE privilege on the package for DSNACCQC
- Ownership of the package
- PACKADM authority for the package collection
- SYSADM authority

The owner of the package or plan that contains the CALL statement must also have SELECT authority on the following catalog table spaces:

- SYSIBM.SYSINDEXES
- SYSIBM.SYSINDEXPART
- SYSIBM.SYSPACKDEP
- SYSIBM.SYSPLANDEP
- SYSIBM.SYSTABLEPART
- SYSIBM.SYSTABLES
- SYSIBM.SYSTABLESPACE

DSNACCQC syntax diagram

The following syntax diagram shows the SQL CALL statement for invoking DSNACCQC. Because the linkage convention for DSNACCQC is GENERAL, you cannot pass null values for input parameters. For character parameters that you are not using, specify an empty string ('').

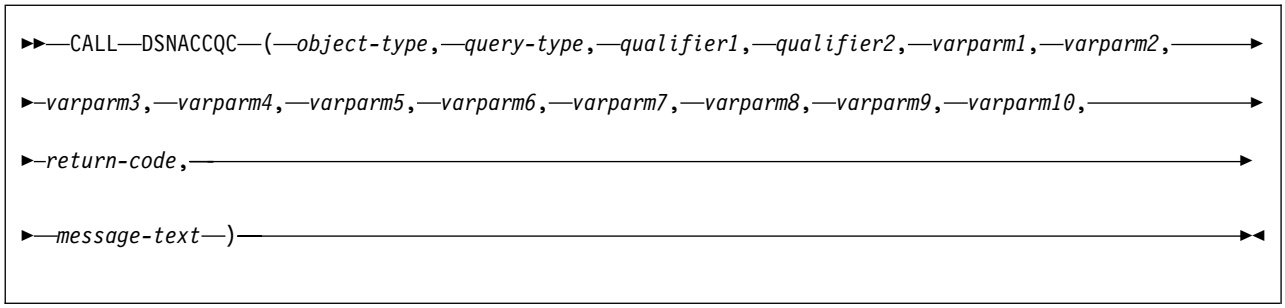

DSNACCQC option descriptions

object-type
Specifies whether you want information about indexes or table spaces. This is an input parameter of type INTEGER. The contents must be one of the following values:

0 Obtain information about indexes

1 Obtain information about table spaces

query-type
Specifies the type of information that you want to obtain. This is an input parameter of type INTEGER.

If *object-type* is 0 (indexes), *query-type* must be one of the following values:

0 Obtains information about indexes on which RUNSTATS needs to be run.

1 Obtains information about indexes with clustering problems.

2 Obtains information about indexes with more than a user-specified number of levels.

3 Obtains information about indexes with a LEAFDIST value that is higher than a user-specified percentage.

4 Obtains information about type 1 indexes.

5 Obtains information about nonunique indexes with long RID chains.

6 Obtains information about indexes that are not used in static SQL statements.

7 Obtains information about indexes on which the STOSPACE utility has not been run.

8 Obtains information about indexes that have exceeded their allocated primary space quantity.

If *object-type* is 1 (table spaces), *query-type* must be one of the following values:

0 Obtains information about table spaces on which RUNSTATS needs to be run.

1 Obtains information about table spaces with more than a user-specified percentage of relocated rows.

2 Obtains information about table spaces with more than a user-specified percentage of space that is occupied by dropped tables.

3 Obtains information about table spaces with table space locking.

4 Obtains information about simple table spaces with more than one table.

5 Obtains information about table spaces on which the STOSPACE utility has not been run.

6 Obtains information about table spaces that have exceeded their allocated primary space quantity.

qualifier1
Narrows the search for objects that match *query-type* to a specified set of database names. This is an input parameter of type VARCHAR(255).

The format of this parameter is the same as the format of *pattern-expression* in an SQL LIKE predicate. *pattern-expression* is described in Chapter 4 of *DB2 SQL Reference*.

For example, to obtain information about table spaces or indexes in all databases with names that begin with ACCOUNT, specify this value for *qualifier1:*

```
ACCOUNT%
```

qualifier2
Narrows the search for objects that match *query-type* to a specified set of creator names. A creator name is the value of column CREATOR in SYSIBM.SYSTABLESPACE for table space queries, or SYSIBM.SYSINDEXES for index queries. This is an input parameter of type VARCHAR(255).

The format of this parameter is the same as the format of *pattern-expression* in an SQL LIKE predicate. *pattern-expression* is described in Chapter 4 of *DB2 SQL Reference*.

For example, to obtain information about table spaces or indexes with creators that begin with DSN8, specify this value for *qualifier2:*

```
DSN8%
```

varparm1, *varparm2*, *varparm3*
The meanings of these parameters vary with *object-type* and *query-type*. See Table 115 on page 645 for the meaning of each parameter for table space queries. See Table 116 on page 645 for the meaning of each parameter for index queries.

These are input parameters of type VARCHAR(255).

varparm4 **through** *varparm10*
These variables are reserved for future use. Specify an empty string ('') for each parameter value.

These are input parameters of type VARCHAR(255).

return-code
Specifies the return code from the DSNACCQC call, which is one of the following values:

0 DSNACCQC executed successfully.

12 An error occurred during DSNACCQC execution.

return-code is an output parameter of type INTEGER.

message-text
If an error occurs while DSNACCQC executes, contains information about the error. If the error is an SQL error, *message-text* also contains the formatted SQLCA. The message text consists of from one to eleven lines, each with a length of 121 bytes. The last byte of each line is a new-line character. *message-text* is an output parameter of type VARCHAR(1331).

Table 115. Variable input parameter values for DSNACCQC table space queries

query-type	**Parameter**	**Value**
0	*varparm1*	Timestamp in character format (*yyyy-mm-dd-hh.mm.ss.nnnnnn*). DSNACCQC returns information about table spaces on which RUNSTATS was run before this time or was never run.
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
1	*varparm1*	Character representation of a number between 0 and 100, which indicates the maximum acceptable percentage of relocated table rows. DSNACCQC returns information about table spaces for which (((FARINDREF+NEARINDREF)/CARD)*100)>*varparm1*.
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
2	*varparm1*	Character representation of a number between 0 and 100, which indicates the maximum acceptable percentage space that is occupied by rows of dropped tables. DSNACCQC returns information about table spaces for which PERCDROP>*varparm1*.
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
3	*varparm1*	Not used. Specify an empty string ('').
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
4	*varparm1*	Not used. Specify an empty string ('').
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
5	*varparm1*	Not used. Specify an empty string ('').
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
6	*varparm1*	Not used. Specify an empty string ('').
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').

Table 116. Variable input parameter values for DSNACCQC index queries

query-type	**Parameter**	**Value**
0	*varparm1*	Timestamp in character format (*yyyy-mm-dd-hh.mm.ss.nnnnnn*). DSNACCQC returns information about indexes on which RUNSTATS was run before this time or was never run.
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').

Table 116. Variable input parameter values for DSNACCQC index queries (continued)

query-type	**Parameter**	**Value**
1	*varparm1*	Character representation of a number between 0 and 100, which indicates the maximum acceptable percentage of table rows that are far from their optimal position. DSNACCQC returns information about indexes for which ((FAROFFPOSF/CARDF)*100)>*varparm1*.
	varparm2	Character representation of a number between 0 and 100, which indicates the maximum acceptable percentage of table rows that are near but not at their optimal position. DSNACCQC returns information about indexes for which ((NEAROFFPOSF/CARDF)*100)>*varparm2*.
	varparm3	Character representation of a number between 0 and 100, which indicates the minimum acceptable percentage of table rows that are in clustering order. DSNACCQC returns information about indexes for which CLUSTERRATIO<*varparm3*.
2	*varparm1*	Character representation of a number that indicates the maximum acceptable number of index levels. DSNACCQC returns information about indexes for which NLEVELS>*varparm1*.
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
3	*varparm1*	Character representation of a number that indicates the maximum acceptable value for 100 times the average number of leaf pages between successive active leaf pages of the index. DSNACCQC returns information about indexes for which LEAFDIST>*varparm1*.
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
4	*varparm1*	Not used. Specify an empty string ('').
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
5	*varparm1*	Character representation of a number that indicates the maximum acceptable average length for RID chains. DSNACCQC returns information about indexes for which ((CARDF*1.0)/FULLKEYCARDF)>*varparm1*.
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
6	*varparm1*	Not used. Specify an empty string ('').
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
7	*varparm1*	Not used. Specify an empty string ('').
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').
8	*varparm1*	Not used. Specify an empty string ('').
	varparm2	Not used. Specify an empty string ('').
	varparm3	Not used. Specify an empty string ('').

Example of DSNACCQC invocation

Suppose that you want information on indexes on which RUNSTATS has never been run. You want information only on indexes in databases whose names begin with DSNCC. The parameter declarations and DSNACCQC call looks like this:

```
DCL OBJTYPE    FIXED BIN(31);
DCL QUERY      FIXED BIN(31);
DCL DBQUAL     CHAR(255) VARYING;
DCL CREATQUAL  CHAR(255) VARYING;
DCL VARPARM1   CHAR(255) VARYING;
DCL VARPARM2   CHAR(255) VARYING;
DCL VARPARM3   CHAR(255) VARYING;
DCL VARPARM4   CHAR(255) VARYING;
DCL VARPARM5   CHAR(255) VARYING;
DCL VARPARM6   CHAR(255) VARYING;
DCL VARPARM7   CHAR(255) VARYING;
DCL VARPARM8   CHAR(255) VARYING;
DCL VARPARM9   CHAR(255) VARYING;
DCL VARPARM10  CHAR(255) VARYING;
DCL RC         FIXED BIN(31);
DCL MSGTEXT    CHAR(1331) VARYING;
DCL IXTABLE SQL TYPE IS RESULT_SET_LOCATOR VARYING;

OBJTYPE=0;
QUERY=0;
DBQUAL='DSNCC%';
CREATQUAL='%';
VARPARM1='0001-01-01-00.00.00.000000';
VARPARM2='';
VARPARM3='';
VARPARM4='';
VARPARM5='';
VARPARM6='';
VARPARM7='';
VARPARM8='';
VARPARM9='';
VARPARM10='';

EXEC SQL CALL SYSPROC.DSNACCQC(:OBJTYPE, :QUERY, :DBQUAL,
 :CREATQUAL, :VARPARM1, :VARPARM2,  :VARPARM3,
 :VARPARM4,  :VARPARM5, :VARPARM6,  :VARPARM7,
 :VARPARM8,  :VARPARM9, :VARPARM10,
 :RC, :MSGTEXT);
```

DSNACCQC output

In addition to the output parameters described in “DSNACCQC option descriptions” on page 643, DSNACCQC returns one result set. The format of the result set varies, depending on whether you are retrieving index information (*object-type*=0) or table space information (*object-type*=1). Table 117 shows the columns of a result set row and the DB2 catalog table that is the source of information for each column for table space queries. Table 118 on page 648 shows the columns of a result set row and the DB2 catalog table that is the source of information for index queries.

Table 117. Result set columns for DSNACCQC table space queries

Column name	Data type	DB2 catalog table that is the data source
NAME	CHAR(8)	SYSTABLESPACE
CREATOR	CHAR(8)	SYSTABLESPACE
BPOOL	CHAR(8)	SYSTABLESPACE
LOCKRULE	CHAR(1)	SYSTABLESPACE

Table 117. Result set columns for DSNACCQC table space queries (continued)

Column name	Data type	DB2 catalog table that is the data source
LOCKMAX	INTEGER	SYSTABLESPACE
CLOSERULE	CHAR(1)	SYSTABLESPACE
ENCODING_SCHEME	CHAR(1)	SYSTABLESPACE
LOCKPART	CHAR(1)	SYSTABLESPACE
MAXROWS	SMALLINT	SYSTABLESPACE
PARTITIONS	SMALLINT	SYSTABLESPACE
TYPE	CHAR(1)	SYSTABLESPACE
SEGSIZE	SMALLINT	SYSTABLESPACE
SPACE	INTEGER	SYSTABLESPACE
NTABLES	SMALLINT	SYSTABLESPACE
STATUS	CHAR(1)	SYSTABLESPACE
STATSTIME	TIMESTAMP	SYSTABLESPACE
ERASERULE	CHAR(1)	SYSTABLESPACE
DBNAME	CHAR(8)	SYSTABLESPACE
DSETPASS	CHAR(8)	SYSTABLESPACE
LOG	CHAR(1)	SYSTABLESPACE
DSSIZE	INTEGER	SYSTABLESPACE
SBCS_CCSID	INTEGER	SYSTABLESPACE

Table 118. Result set columns for DSNACCQC index queries

Column name	Data type	DB2 catalog table that is the data source
CREATOR	CHAR(8)	SYSINDEXES
NAME	VARCHAR(18)	SYSINDEXES
TBCREATOR	CHAR(8)	SYSINDEXES
TBNAME	VARCHAR(18)	SYSINDEXES
UNIQUERULE	CHAR(1)	SYSINDEXES
INDEXTYPE	CHAR(1)	SYSINDEXES
INDEXSPACE	CHAR(8)	SYSINDEXES
CLUSTERING	CHAR(1)	SYSINDEXES
ERASERULE	CHAR(1)	SYSINDEXES
CLOSERULE	CHAR(1)	SYSINDEXES
COLCOUNT	SMALLINT	SYSINDEXES
DBID	SMALLINT	SYSINDEXES
DBNAME	CHAR(8)	SYSINDEXES
BPOOL	CHAR(8)	SYSINDEXES
PGSIZE	SMALLINT	SYSINDEXES
DSETPASS	CHAR(8)	SYSINDEXES
PIECESIZE	INTEGER	SYSINDEXES
COPY	CHAR(1)	SYSINDEXES

Table 118. Result set columns for DSNACCQC index queries (continued)

Column name	Data type	DB2 catalog table that is the data source
PARTITION_COUNT	INTEGER	SYSINDEXPART[1]

Notes:

1. The value of PARTITION_COUNT is COUNT(DISTINCT PARTITION) for partitioning indexes or 0 for nonpartitioning indexes. The PARTITION column is in SYSIBM.SYSINDEXPART.

To obtain the information from the result set, you can write your client program to retrieve information from a one result set with known contents. However, for greater flexibility, you might want to write your client program to retrieve data from an unknown number of result sets with unknown contents. Both techniques are shown in Part 6 of *DB2 Application Programming and SQL Guide*.

The DB2 UDB Control Center partition information stored procedure (DSNACCAV)

The information under this heading is Product-sensitive Programming Interface and Associated Guidance Information, as defined in “Notices” on page 725.

DSNACCAV is a sample stored procedure that gives you information about partitions in your table spaces and indexes. You can use DSNACCAV to obtain the following types of information:

- Partitions that need to be image copied
- Partitions that are in a restricted state
- Partitions on which RUNSTATS needs to be run
- Partitions on which REORG needs to be run
- Partitions that exceed a user-specified number of extents

Environment

DSNACCAV runs in a WLM-established stored procedures address space.

Authorization required

To execute the CALL statement, the owner of the package or plan that contains the CALL statement must have one or more of the following privileges on each package that the stored procedure uses:

- The EXECUTE privilege on the package for DSNACCAV
- Ownership of the package
- PACKADM authority for the package collection
- SYSADM authority

The owner of the package or plan that contains the CALL statement must also have SELECT authority on the following catalog table spaces:

- SYSIBM.SYSCOPY
- SYSIBM.SYSINDEXES
- SYSIBM.SYSINDEXPART
- SYSIBM.SYSTABLEPART
- SYSIBM.SYSTABLES
- SYSIBM.SYSTABLESPACE

DSNACCAV syntax diagram

The following syntax diagram shows the SQL CALL statement for invoking DSNACCAV. Because the linkage convention for DSNACCAV is GENERAL WITH NULLS, if you pass parameters in host variables, you need to include a null indicator with every host variable. Null indicators for input host variables must be initialized before you execute the CALL statement.

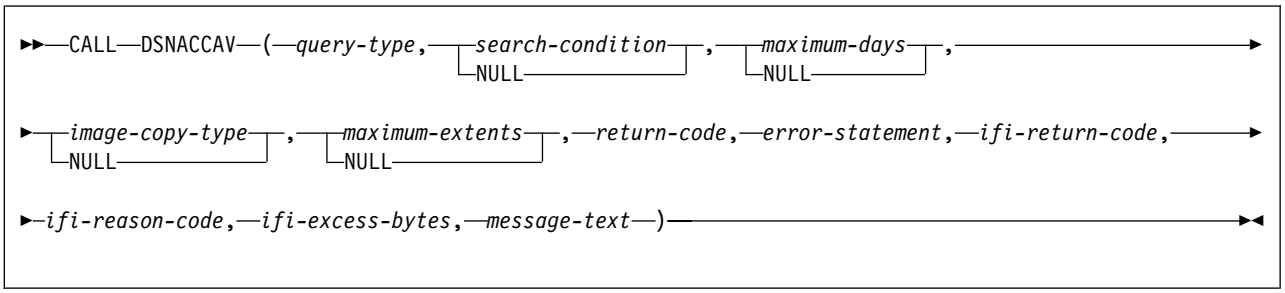

DSNACCAV option descriptions

query-type

Specifies the type of information that you want to obtain. This is an input parameter of type VARCHAR(20). The contents must be one of the following values:

COPY TABLESPACE
Obtains information about table space partitions for which image copies need to be made.

COPY INDEX
Obtains information about index partitions for which image copies need to be made.

RESTRICT TABLESPACE
Obtains information about table space partitions that are in a restricted status.

RESTRICT INDEX
Obtains information about index partitions that are in a restricted status.

RUNSTATS TABLESPACE
Obtains information about table space partitions on which RUNSTATS needs to be run.

RUNSTATS INDEX
Obtains information about index partitions on which RUNSTATS needs to be run.

REORG TABLESPACE
Obtains information about table space partitions on which REORG needs to be run.

REORG INDEX
Obtains information about index partitions on which REORG needs to be run.

EXTENTS TABLESPACE
Obtains information about table space partitions that have used more than a user-specified number of extents.

EXTENTS INDEX
Obtains information about index partitions that have used more than a user-specified number of extents.

search-condition
Narrows the search for objects that match *query-type.* This is an input parameter of type VARCHAR(4096).

The format of this parameter is the same as the format of *search-condition* in an SQL *where-clause. search-condition* is described in Chapter 4 of *DB2 SQL Reference*.

If the call is executed to obtain table space information, *search-condition* can include any column in SYSIBM.SYSTABLESPACE. If the call is executed to obtain index information, *search-condition* can include any column in SYSIBM.SYSINDEXES. Each column name must be preceded by the string 'A.'.

For example, to obtain information about table spaces with creator ADMF001, specify this value for *search-condition:*

```
A.CREATOR='ADMF001'
```

maximum-days
Specifies the maximum number of days that should elapse between executions of the REORG, RUNSTATS, or COPY utility. DSNACCAV uses this value as the criterion for determining which table space or index partitions need to have the utility that you specified in *query-type* run against them. This value can be specified if *query-type* has one of the following values:
- COPY TABLESPACE
- COPY INDEX
- RUNSTATS TABLESPACE
- RUNSTATS INDEX
- REORG TABLESPACE
- REORG INDEX

maximum-days is an input parameter of type INTEGER.

image-copy-type
Specifies the types of image copies about which DSNACCAV should give you information. This value can be specified if *query-type* is COPY TABLESPACE or COPY INDEX. *image-copy-type* is an input parameter of type CHAR(1). The contents must be one of the following values:

B Specifies that you want information about partitions for which the most recent image copy was either a full image copy or an incremental image copy

F Specifies that you want information about partitions for which the most recent image copy was a full image copy

I Specifies that you want information about partitions for which the most recent image copy was an incremental image copy

maximum-extents
Specifies the maximum number of extents that a table space or index partition should use. This value can be specified if *query-type* is one of the following values:
- REORG TABLESPACE
- REORG INDEX
- EXTENTS TABLESPACE
- EXTENTS INDEX

maximum-extents is an input parameter of type INTEGER.

return-code
Specifies the return code from the DSNACCAV call, which is one of the following values:

0 DSNCCAV executed successfully.

12 An error occurred during DSNCCAV execution.

return-code is an output parameter of type INTEGER.

error-statement
If *return-code* is not 0, specifies the SQL statement or DB2 command that DB2 was executing when the error occurred. *error-statement* is an output parameter of type VARCHAR(8012).

ifi-return-code
When *query-type* is RESTRICT TABLESPACE, RESTRICT INDEX, COPY TABLESPACE, or REORG TABLESPACE, specifies the return code from the IFI call that submitted a DISPLAY DATABASE command to obtain information about restricted objects. *ifi-return-code* is an output parameter of type INTEGER.

ifi-reason-code
When *query-type* is RESTRICT TABLESPACE, RESTRICT INDEX, COPY TABLESPACE, or REORG TABLESPACE, specifies the reason code from the IFI call that submitted a DISPLAY DATABASE command to obtain information about restricted objects. *ifi-reason-code* is an output parameter of type INTEGER.

ifi-excess-bytes
When *query-type* is RESTRICT TABLESPACE, RESTRICT INDEX, COPY TABLESPACE, or REORG TABLESPACE, specifies the number of bytes that did not fit in the return area for the IFI call that submitted a DISPLAY DATABASE command to obtain information about restricted objects. *ifi-excess-bytes* is an output parameter of type INTEGER.

message-text
If an SQL error occurs while DSNACCAV executes, contains information about the error, including the formatted SQLCA. The message text consists of from one to ten lines, each with a length of 121 bytes. The last byte of each line is a new-line character. *message-text* is an output parameter of type VARCHAR(1210).

Example of DSNACCAV invocation

Suppose that you want information on table space partitions that are in a restricted status. You want information only on table spaces that are in databases whose names begin with DSNCC. The parameter declarations and DSNACCAV call looks like this:

```
DCL QUERY      CHAR(20) VARYING;
DCL CRITERIA   CHAR(4096) VARYING;
DCL NUMDAYS    FIXED BIN(31);
DCL OPTYPE     CHAR(1) VARYING;
DCL EXTENTS    FIXED BIN(31);
DCL RC         FIXED BIN(31);
DCL STMT       CHAR(8012) VARYING;
DCL IFIRC      FIXED BIN(31);
DCL IFIREASON  FIXED BIN(31);
DCL IFIEXCESS  FIXED BIN(31);
DCL STMT       CHAR(8012) VARYING;
DCL MSGTEXT    CHAR(1331) VARYING;
DCL IND1 FIXED BIN(15);
```

```
DCL IND2 FIXED BIN(15);
DCL IND3 FIXED BIN(15);
DCL IND4 FIXED BIN(15);
DCL IND5 FIXED BIN(15);
DCL IND6 FIXED BIN(15);
DCL IND7 FIXED BIN(15);
DCL IND8 FIXED BIN(15);
DCL IND9 FIXED BIN(15);
DCL IND10 FIXED BIN(15);
DCL IND11 FIXED BIN(15);
DCL CMDMSG  SQL TYPE IS RESULT_SET_LOCATOR VARYING;
DCL TSTABLE SQL TYPE IS RESULT_SET_LOCATOR VARYING;

QUERY='RESTRICT TABLESPACE';
IND1=0;
Criteria='A.DBNAME LIKE ''DSNCC%''';
IND2=0;
numdays=0;
IND3=0;
optype='';
IND4=0;
extents=0;
IND5=0;

EXEC SQL CALL SYSPROC.DSNACCAV(:QUERY :IND1, :CRITERIA :IND2,
 :NUMDAYS :IND3, :OPTYPE :IND4, :EXTENTS :IND5,
 :RC :IND6, :STMT :IND7, :IFIRC :IND8,
 :IFIREASON :IND9, :IFIEXCESS :IND10, :MSGTEXT :IND11);
```

DSNACCAV output

In addition to the output parameters described in “DSNACCAV option descriptions” on page 650, DSNACCAV returns two result sets.

The first result set contains the text from commands that DB2 executes, formatted into 80-byte records. Table 119 shows the format of the first result set.

The second result set contains partition information. The format of the second result set varies, depending on whether you request table space or index information. Table 120 shows the columns of a result set row and the DB2 catalog table that is the source of information for each column for table space queries. Table 122 on page 655 shows the same information for index queries.

The number of rows that are returned in the second result set varies with *query-type*. Table 124 on page 657 shows the number and types of rows that are returned from an invocation of DSNACCAV for each *query-type*.

Table 119. Result set row for DSNACCAV command output

Column name	Data type	Contents
RS_SEQUENCE	INTEGER	Sequence number of the output line
RS_DATA	CHAR(80)	A line of command output

Table 120. Result set row for DSNACCAV table space queries

Column name	Data type	DB2 catalog table that is the data source
NAME	CHAR(8)	SYSTABLESPACE
CREATOR	CHAR(8)	SYSTABLESPACE
BPOOL	CHAR(8)	SYSTABLESPACE

Table 120. Result set row for DSNACCAV table space queries (continued)

Column name	Data type	DB2 catalog table that is the data source
LOCKRULE	CHAR(1)	SYSTABLESPACE
LOCKMAX	INTEGER	SYSTABLESPACE
CLOSERULE	CHAR(1)	SYSTABLESPACE
ENCODING_SCHEME	CHAR(1)	SYSTABLESPACE
LOCKPART	CHAR(1)	SYSTABLESPACE
MAXROWS	SMALLINT	SYSTABLESPACE
PARTITIONS	SMALLINT	SYSTABLESPACE
TYPE	CHAR(1)	SYSTABLESPACE
SEGSIZE	SMALLINT	SYSTABLESPACE
SPACE	INTEGER	SYSTABLESPACE
NTABLES	SMALLINT	SYSTABLESPACE
STATUS	CHAR(1)	SYSTABLESPACE
STATSTIME	TIMESTAMP	SYSTABLESPACE
ERASERULE	CHAR(1)	SYSTABLESPACE
DBNAME	CHAR(8)	SYSTABLESPACE
DSETPASS	CHAR(8)	SYSTABLESPACE
LOG	CHAR(1)	SYSTABLESPACE
DSSIZE	INTEGER	SYSTABLESPACE
PARTITION	SMALLINT	SYSTABLEPART
OPERATIONTIME	TIMESTAMP	SYSCOPY or SYSTABLEPART[1]
DAYS	INTEGER	SYSCOPY or SYSTABLEPART[2]
PERCOFFPOS	SMALLINT	SYSINDEXPART[3]
PERCINDREF	SMALLINT	SYSTABLEPART[4]
PERCDROP	SMALLINT	SYSTABLEPART
EXTENTS	INTEGER	None[5]
REASON	CHAR(18)	None[6]
SBCS_CCSID	INTEGER	SYSTABLESPACE

Notes:

1. If *query-type* is COPY TABLESPACE or REORG TABLESPACE, the value of OPERATIONTIME is the value of the TIMESTAMP column in SYSIBM.SYSCOPY.

   If *query-type* is RUNSTATS TABLESPACE, the value of OPERATIONTIME is the value of the TIMESTAMP column in SYSIBM.SYSTABLEPART.
2. DAYS is the number of days since the last invocation of the utility. This column is derived from the OPERATIONTIME column.
3. PERCOFFPOS=(NEAROFFPOSF+FAROFFPOSF)*100/CARDF
4. PERCINDREF=(NEARINDREF+FARINDREF)*100/CARD
5. EXTENTS is the number of data set extents that the partition is using.
6. REASON is the reason that the row is in the result set. See Table 121 on page 655 for values of REASON for each value of *query-type*.

Table 121. Values of the REASON result set column for table space queries

query-type	**REASON value**	**REASON meaning**
COPY TABLESPACE	Status from DISPLAY DATABASE command output	Table space is in restricted status COPY
	DAYS	Days since last COPY exceeds *maximum-days* value
RESTRICT TABLESPACE	Status from DISPLAY DATABASE command output	Table space is in restricted status
RUNSTATS TABLESPACE	LOAD	LOAD was run after RUNSTATS
	REORG	REORG was run after RUNSTATS
	RECOVER	RECOVER was run after RUNSTATS
	DAYS	Days since last RUNSTATS exceeds *maximum-days* value

REORG TABLESPACE	LIMIT	One of the following reasons:
		• A clustering index meets this condition:
		`((NEAROFFPOSF+FAROFFPOSF)*100/CARDF)>10`
		• A partition meets either of these conditions:
		`((NEARINDREF+FARINDREF)*100/CARD)>10`
		`PERCDROP>10`

	DAYS	Days since last REORG exceeds *maximum-days* value
	EXTENTS	Number of partition extents exceeds *maximum-extents* value
	Status from DISPLAY DATABASE command output	The table space is in restricted status REORP
EXTENTS TABLESPACE	EXTENTS	Number of partition extents exceeds *maximum-extents* value

Table 122. Result set row for DSNACCAV index queries

Column name	**Data type**	**DB2 catalog table that is the data source**
CREATOR	CHAR(8)	SYSINDEXES
NAME	VARCHAR(18)	SYSINDEXES
TBCREATOR	CHAR(8)	SYSINDEXES
TBNAME	VARCHAR(18)	SYSINDEXES
UNIQUERULE	CHAR(1)	SYSINDEXES
INDEXTYPE	CHAR(1)	SYSINDEXES
INDEXSPACE	CHAR(8)	SYSINDEXES
CLUSTERING	CHAR(1)	SYSINDEXES
ERASERULE	CHAR(1)	SYSINDEXES
CLOSERULE	CHAR(1)	SYSINDEXES

Table 122. Result set row for DSNACCAV index queries (continued)

Column name	Data type	DB2 catalog table that is the data source
COLCOUNT	SMALLINT	SYSINDEXES
DBID	SMALLINT	SYSINDEXES
DBNAME	CHAR(8)	SYSINDEXES
BPOOL	CHAR(8)	SYSINDEXES
PGSIZE	SMALLINT	SYSINDEXES
DSETPASS	CHAR(8)	SYSINDEXES
PIECESIZE	INTEGER	SYSINDEXES
COPY	CHAR(1)	SYSINDEXES
PARTITIONS	SMALLINT	SYSINDEXPART[1]
PARTITION	SMALLINT	SYSINDEXPART
OPERATIONTIME	TIMESTAMP	SYSCOPY or SYSINDEXPART[2]
DAYS	INTEGER	SYSCOPY or SYSTABLEPART[3]
LEAFDIST	INTEGER	SYSINDEXPART
EXTENTS	INTEGER	None[4]
REASON	CHAR(18)	None[5]

Notes:

1. PARTITIONS is derived from the PARTITION column through this SELECT statement:

```
SELECT IXNAME,IXCREATOR,MAX(PARTITION) AS PARTITIONS
  FROM SYSIBM.SYSINDEXPART
  GROUP BY IXNAME,IXCREATOR;
```

2. If *query-type* is COPY INDEX or REORG INDEX, the value of OPERATIONTIME is the value of the TIMESTAMP column in SYSIBM.SYSCOPY.

   If *query-type* is RUNSTATS INDEX, the value of OPERATIONTIME is the value of the TIMESTAMP column in SYSIBM.SYSINDEXPART.

3. DAYS is the number of days since the last invocation of the utility. This column is derived from the OPERATIONTIME column.
4. EXTENTS is the number of data set extents that the partition is using.
5. REASON is the reason that the row is in the result set. See Table 123 for values of REASON for each value of *query-type*.

Table 123. Values of the REASON result set column for index queries

query-type	REASON value	REASON meaning
COPY INDEX	LIMIT	Index is in restricted status ICOPY
	DAYS	Days since last COPY exceeds *maximum-days* value
RESTRICT INDEX	Status from DISPLAY DATABASE command output	Index is in restricted status

Table 123. Values of the REASON result set column for index queries (continued)

query-type	**REASON value**	**REASON meaning**
RUNSTATS INDEX	TABLESPACE LOAD	LOAD was run on the associated table space after RUNSTATS
	TABLESPACE REORG	REORG was run on the associated table space after RUNSTATS
	TABLESPACE RECOVER	RECOVER was run on the associated table space after RUNSTATS
	REBUILT	REBUILD was run after RUNSTATS
	DAYS	Days since last RUNSTATS exceeds *maximum-days* value
REORG INDEX	LIMIT	LEAFDIST exceeds the recommended limit of 200
	DAYS	Days since last REORG exceeds *maximum-days* value
	EXTENTS	Number of partition extents exceeds *maximum-extents* value
EXTENTS INDEX	EXTENTS	Number of partition extents exceeds *maximum-extents* value

Table 124. Rows of the second DSNACCAV result set for each query type

query-type	**Rows returned**
COPY TABLESPACE	One row for: • Each tablespace partition that has not been copied within the number of days specified by the *maximum-days* parameter • The most recent copy of each data set in a nonpartitioned tablespace • Each table space partition that is in COPY-pending status
COPY INDEX	One row for: • Each index space partition that has not been copied within the number of days specified by the *maximum-days* parameter • The most recent copy of each data set in a nonpartitioned index space • Each index space partition that is in ICOPY-pending status
RESTRICT TABLESPACE	One row for each table space in the subsystem that meets the criteria specified by the *search-criteria* parameter and is in a restricted status
RESTRICT INDEX	One row for each index in the subsystem that meets the criteria specified by the *search-criteria* parameter and is in a restricted status
RUNSTATS TABLESPACE	One row for: • Each tablespace partition on which LOAD, REORG, or RECOVER was run after the last time RUNSTATS was run • Each table space partition on which RUNSTATS was not run within the number of days specified by the *maximum-days* parameter

Table 124. Rows of the second DSNACCAV result set for each query type (continued)

query-type	**Rows returned**
RUNSTATS INDEX	One row for: • Each index partition that is defined on a table on which LOAD, REORG, or RECOVER was run after the last time RUNSTATS was run • Each index partition on which RUNSTATS was not run within the number of days specified by the *maximum-days* parameter • Each index partition on which REBUILD was run after the last time RUNSTATS was run
REORG TABLESPACE	One row for: • Each table space partition that is in REORG-pending status • Each table space partition for which the number of data set extents exceeds the value specified by *maximum-extents* • Each table space partition for which the clustering index that is associated with the table has (NEAROFFPOSF+FAROFFPOSF)*100/CARDF>10 • Each table space partition for which (NEARINDREF+FARINDREF)*100/CARD>10 • Each table space partition for which PERCDROP>10 • Each table space partition on which REORG was not run within the number of days specified by the *maximum-days* parameter
REORG INDEX	One row for: • Each index partition for which LEAFDIST>200 • Each index partition for which the number of data set extents exceeds the value specified by *maximum-extents* • Each index partition on which REORG was not run within the number of days specified by the *maximum-days* parameter
EXTENTS TABLESPACE	One row for each table space partition for which the number of data set extents exceeds the value specified by *maximum-extents*
EXTENTS INDEX	One row for each index partition for which the number of data set extents exceeds the value specified by *maximum-extents*

To obtain the information from the result sets, you can write your client program to retrieve information from two result sets with known contents. However, for greater flexibility, you might want to write your client program to retrieve data from an unknown number of result sets with unknown contents. Both techniques are shown in Part 6 of *DB2 Application Programming and SQL Guide*.

The DB2 real-time statistics stored procedure

The information under this heading is Product-sensitive Programming Interface and Associated Guidance Information, as defined in “Notices” on page 725.

The DSNACCOR stored procedure is a sample stored procedure that makes recommendations to help you maintain your DB2 databases. In particular, DSNACCOR performs these actions:

- Recommends when you should reorganize, image copy, or update statistics for table spaces or index spaces
- Indicates table spaces or index spaces that have exceeded their data set extents

- Indicates whether an object for which it recommends an action is in a restricted state when DSNACCOR runs

DSNACCOR uses data from the SYSIBM.TABLESPACESTATS and SYSIBM.INDEXSPACESTATS real-time statistics tables to make its recommendations. DSNACCOR provides its recommendations in a result set.

DSNACCOR uses the set of criteria shown in "Formulas for recommending actions" on page 667 to evaluate table spaces and index spaces. By default, DSNACCOR evaluates *all* table spaces and index spaces in the subsystem that have entries in the real-time statistics tables. However, you can override this default through input parameters.

Important information about DSNACCOR recommendations:

- DSNACCOR makes recommendations based on general formulas that require input from the user about the maintenance policies for a subsystem. These recommendations might not be accurate for every installation.
- If the real-time statistics tables contain information for only a small percentage of your DB2 subsystem, the recommendations that DSNACCOR makes might not be accurate for the entire subsystem.
- Before you perform any action that DSNACCOR recommends, ensure that the object for which DSNACCOR makes the recommendation is available, and that it is possible to perform the recommended action on that object. For example, before you can perform an image copy on an index, the index must have the COPY YES attribute.

Environment

DSNACCOR must run in a WLM-established stored procedure address space.

DSNACCOR creates and uses declared temporary tables. Therefore, before you can invoke DSNACCOR, you need to create a TEMP database and segmented table spaces in the TEMP database. Specify a 4KB buffer pool when you create the TEMP database. For information on creating TEMP databases and table spaces, see the CREATE DATABASE and CREATE TABLESPACE sections in Chapter 5 of *DB2 SQL Reference*.

Before you can invoke DSNACCOR, the real-time statistics tables, SYSIBM.TABLESPACESTATS and SYSIBM.INDEXSPACESTATS, must exist, and the real-time statistics database must be started. See "Appendix F. Real-time statistics tables" on page 699 for information on the real-time statistics tables.

Authorization required

To execute the CALL DSNACCOR statement, the owner of the package or plan that contains the CALL statement must have one or more of the following privileges on each package that the stored procedure uses:

- The EXECUTE privilege on the package for DSNACCOR
- Ownership of the package
- PACKADM authority for the package collection
- SYSADM authority

The owner of the package or plan that contains the CALL statement must also have:

- SELECT authority on the real-time statistics tables

- The DISPLAY system privilege

DSNACCOR syntax diagram

The following syntax diagram shows the CALL statement for invoking DSNACCOR. Because the linkage convention for DSNACCOR is GENERAL WITH NULLS, if you pass parameters in host variables, you need to include a null indicator with every host variable. Null indicators for input host variables must be initialized before you execute the CALL statement.

DSNACCOR option descriptions

In the following option descriptions, the default value for an input parameter is the value that DSNACCOR uses if you specify a null value.

QueryType
Specifies the types of actions that DSNACCOR recommends. This field contains one or more of the following values. Each value is enclosed in single quotation marks and separated from other values by a space.

ALL Make recommendations for all of the following actions.

COPY Make a recommendation on whether to perform an image copy.

RUNSTATS Make a recommendation on whether to perform RUNSTATS.

REORG Make a recommendation on whether to perform REORG. Choosing this value causes DSNACCOR to process the EXTENTS value also.

EXTENTS Indicate when data sets have exceeded a user-specified extents limit.

RESTRICT Indicate which objects are in a restricted state.

This is an input parameter of type VARCHAR(40). The default is 'ALL'.

ObjectType
Specifies the types of objects for which DSNACCOR recommends actions:

ALL Table spaces and index spaces.

TS Table spaces only.

IX Index spaces only.

This is an input parameter of type VARCHAR(3). The default is 'ALL'.

ICType
Specifies the types of image copies for which DSNACCOR should make recommendations:

F Full image copy.

I Incremental image copy. This value is valid for table spaces only.

B Full image copy or incremental image copy.

This is an input parameter of type VARCHAR(1). The default is 'B'.

StatsSchema
Specifies the qualifier for the real-time statistics table names. This is an input parameter of type VARCHAR(128). The default is 'SYSIBM'.

CatlgSchema
Specifies the qualifier for DB2 catalog table names. This is an input parameter of type VARCHAR(128). The default is 'SYSIBM'.

LocalSchema
Specifies the qualifier for the names of tables that DSNACCOR creates. This is an input parameter of type VARCHAR(128). The default is 'DSNACC'.

ChkLvl
Specifies the types of checking that DSNACCOR performs, and indicates whether to include objects that fail those checks in the DSNACCOR recommendations result set. This value is the sum of any combination of the following values:

0 DSNACCOR performs none of the following actions.

1 For objects that are listed in the recommendations result set, check the SYSTABLESPACE or SYSINDEXES catalog tables to ensure that those objects have not been deleted. If value 16 is *not* also chosen, exclude rows for the deleted objects from the recommendations result set.

2 For index spaces that are listed in the recommendations result set, check the SYSTABLES, SYSTABLESPACE, and SYSINDEXES catalog tables to determine the name of the table space that is associated with each index space.

Choosing this value causes DSNACCOR to also check for rows in the recommendations result set for objects that have been deleted but have entries in the real-time statistics tables (value 1). This means that if value 16 is *not* also chosen, rows for deleted objects are excluded from the recommendations result set.

4 Check whether rows that are in the DSNACCOR recommendations result set refer to objects that are in the exception table. For recommendations result set rows that have corresponding exception table rows, copy the contents of the QUERYTYPE column of the exception table to the INEXCEPTTABLE column of the recommendations result set.

8 Check whether objects that have rows in the recommendations result set are restricted. Indicate the restricted status in the OBJECTSTATUS column of the result set.

16 For objects that are listed in the recommendations result set, check the SYSTABLESPACE or SYSINDEXES catalog tables to ensure that those objects have not been deleted (value 1). In result set rows for deleted objects, specify the word ORPHANED in the OBJECTSTATUS column.

32 Exclude rows from the DSNACCOR recommendations result set for index spaces for which the related table spaces have been recommended for REORG. Choosing this value causes DSNACCOR to perform the actions for values 1 and 2.

This is an input parameter of type INTEGER. The default is 7 (values 1+2+4).

Criteria
Narrows the set of objects for which DSNACCOR makes recommendations. This value is the search condition of an SQL WHERE clause. This is an input parameter of type VARCHAR(4096). The default is that DSNACCOR makes recommendations for all table spaces and index spaces in the subsystem.

Unused
A parameter that is reserved for future use. Specify the null value for this parameter. This is an input parameter of type VARCHAR(80).

CRUpdatedPagesPct
Specifies a criterion for recommending a full image copy on a table space or index space. For a table space, if the ratio of distinct updated pages to preformatted pages, expressed as a percentage, is greater than this value, DSNACCOR recommends an image copy. (See item 2 in Figure 44 on page 667.) For an index space, if the ratio of distinct updated pages to preformatted pages, expressed as a percentage, is greater than this value, *and* the number of active pages in the index space or partition is greater than *CRIndexSize*, DSNACCOR recommends an image copy. (See items 2 and 3 in Figure 45 on page 667.) This is an input parameter of type INTEGER. The default is 20.

CRChangesPct
Specifies a criterion for recommending a full image copy on a table space or index space. For a table space, if the ratio of the number INSERTs, UPDATEs, and DELETEs since the last image copy to the total number of rows or LOBs in a table space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends an image copy. (See item 3 in Figure 44 on page 667.) For an index space, if the ratio of the number INSERTs and DELETEs since the last image copy to the total number of entries in the index space or partition, expressed as a percentage, is greater than this value, *and*

the number of active pages in the index space or partition is greater than *CRIndexSize*, DSNACCOR recommends an image copy. (See items 2 and 4 in Figure 45 on page 667.) This is an input parameter of type INTEGER. The default is 10.

CRDaySncLastCopy
Specifies a criterion for recommending a full image copy on a table space or index space. If the number of days since the last image copy is greater than this value, DSNACCOR recommends an image copy. (See item 1 in Figure 44 on page 667 and item 1 in Figure 45 on page 667.) This is an input parameter of type INTEGER. The default is 7.

ICRUpdatedPagesPct
Specifies a criterion for recommending an incremental image copy on a table space. If the ratio of the number of distinct pages updated since the last image copy to the total number of active pages in the table space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends an incremental image copy. (See item 1 in Figure 46 on page 667.) This is an input parameter of type INTEGER. The default is 1.

ICRChangesPct
Specifies a criterion for recommending an incremental image copy on a table space. If the ratio of the number INSERTs, UPDATEs, and DELETEs since the last image copy to the total number of rows or LOBs in a table space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends an incremental image copy. (See item 2 in Figure 46 on page 667.) This is an input parameter of type INTEGER. The default is 1.

CRIndexSize
Combined with *CRUpdatedPagesPct* or *CRChangesPct*, specifies a criterion for recommending a full image copy on an index space. (See items 2, 3, and 4 in Figure 45 on page 667.) This is an input parameter of type INTEGER. The default is 50.

RRTInsDelUpdPct
Specifies a criterion for recommending that the REORG utility should be run on a table space. If the ratio of the sum of INSERTs, UPDATEs, and DELETEs since the last REORG to the total number of rows or LOBs in the table space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends running REORG. (See item 1 in Figure 47 on page 668.) This is an input parameter of type INTEGER. The default is 20.

RRTUnclustInsPct
Specifies a criterion for recommending that the REORG utility should be run on a table space. If the ratio of the number of unclustered INSERTs to the total number of rows or LOBs in the table space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends running REORG. (See item 2 in Figure 47 on page 668.) This is an input parameter of type INTEGER. The default is 10.

RRTDisorgLOBPct
Specifies a criterion for recommending that the REORG utility should be run on a table space. If the ratio of the number of imperfectly chunked LOBs to the total number of rows or LOBs in the table space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends running REORG. (See item 3 in Figure 47 on page 668.) This is an input parameter of type INTEGER. The default is 10.

RRTMassDelLimit
Specifies a criterion for recommending that the REORG utility should be run on

a table space. If the number of mass deletes from a segmented or LOB table space since the last REORG or LOAD REPLACE, or the number of dropped tables from a nonsegmented table space since the last REORG or LOAD REPLACE is greater than this value, DSNACCOR recommends running REORG. (See item 5 in Figure 47 on page 668.) This is an input parameter of type INTEGER. The default is 0.

RRTIndRefLimit
Specifies a criterion for recommending that the REORG utility should be run on a table space. If the ratio of the total number of overflow records that were created since the last REORG or LOAD REPLACE to the total number of rows or LOBs in the table space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends running REORG. (See item 4 in Figure 47 on page 668.) This is an input parameter of type INTEGER. The default is 10.

RRIInsertDeletePct
Specifies a criterion for recommending that the REORG utility should be run on an index space. If the ratio of the sum of the number of index entries that were inserted and deleted since the last REORG to the total number of index entries in the index space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends running REORG. (See item 1 in Figure 48 on page 668.) This is an input parameter of type INTEGER. The default is 20.

RRIAppendInsertPct
Specifies a criterion for recommending that the REORG utility should be run on an index space. If the ratio of the number of index entries that were inserted since the last REORG, REBUILD INDEX, or LOAD REPLACE, and had a key value greater than the maximum key value in the index space or partition, to the number of index entries in the index space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends running REORG. (See item 2 in Figure 48 on page 668.) This is an input parameter of type INTEGER. The default is 10.

RRIPseudoDeletePct
Specifies a criterion for recommending that the REORG utility should be run on an index space. If the ratio of the number of index entries that were pseudo-deleted since the last REORG, REBUILD INDEX, or LOAD REPLACE to the number of index entries in the index space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends running REORG. (See item 3 in Figure 48 on page 668.) This is an input parameter of type INTEGER. The default is 10.

RRIMassDelLimit
Specifies a criterion for recommending that the REORG utility should be run on an index space. If the number of mass deletes from an index space or partition since the last REORG, REBUILD, or LOAD REPLACE is greater than this value, DSNACCOR recommends running REORG. (See item 4 in Figure 48 on page 668.) This is an input parameter of type INTEGER. The default is 0.

RRILeafLimit
Specifies a criterion for recommending that the REORG utility should be run on an index space. If the ratio of the number of index page splits that occurred since the last REORG, REBUILD INDEX, or LOAD REPLACE in which the higher part of the split page was far from the location of the original page, to the total number of active pages in the index space or partition, expressed as a percentage, is greater than this value, DSNACCOR recommends running REORG. (See item 5 in Figure 48 on page 668.) This is an input parameter of type INTEGER. The default is 10.

RRINumLevelsLimit
Specifies a criterion for recommending that the REORG utility should be run on an index space. If the number of levels in the index tree that were added or removed since the last REORG, REBUILD INDEX, or LOAD REPLACE is greater than this value, DSNACCOR recommends running REORG. (See item 6 in Figure 48 on page 668.) This is an input parameter of type INTEGER. The default is 0.

SRTInsDelUpdPct
Combined with *SRTInsDelUpdAbs*, specifies a criterion for recommending that the RUNSTATS utility should be run on a table space. If the ratio of the number INSERTs, UPDATEs, and DELETEs since the last RUNSTATS on a table space or partition, to the total number of rows or LOBs in the table space or partition, expressed as a percentage, is greater than *SRTInsDelUpdPct*, *and* the sum of the number INSERTs, UPDATEs, and DELETEs since the last RUNSTATS on a table space or partition is greater than *SRTInsDelUpdAbs*, DSNACCOR recommends running RUNSTATS. (See items 1 and 2 in Figure 49 on page 668.) This is an input parameter of type INTEGER. The default is 20.

SRTInsDelUpdAbs
Combined with *SRTInsDelUpdAbs*, specifies a criterion for recommending that the RUNSTATS utility should be run on a table space. (See items 1 and 2 in Figure 49 on page 668.) This is an input parameter of type INTEGER. The default is 0.

SRTMassDelLimit
Specifies a criterion for recommending that the RUNSTATS utility should be run on a table space. If the number of mass deletes from a table space or partition since the last REORG or LOAD REPLACE is greater than this value, DSNACCOR recommends running RUNSTATS. (See item 3 in Figure 49 on page 668.) This is an input parameter of type INTEGER. The default is 0.

SRIInsDelUpdPct
Combined with *SRIInsDelUpdAbs*, specifies a criterion for recommending that the RUNSTATS utility should be run on an index space. If the ratio of the number inserted and deleted index entries since the last RUNSTATS on an index space or partition, to the total number of index entries in the index space or partition, expressed as a percentage, is greater than *SRIInsDelUpdPct*, *and* the sum of the number inserted and deleted index entries since the last RUNSTATS on an index space or partition is greater than *SRIInsDelUpdAbs*, DSNACCOR recommends running RUNSTATS. (See items 1 and 2 in Figure 50 on page 668.) This is an input parameter of type INTEGER. The default is 20.

SRIInsDelUpdAbs
Combined with *SRIInsDelUpdPct*, specifies a criterion for recommending that the RUNSTATS utility should be run on an index space. (See items 1 and 2 in Figure 50 on page 668.) This is an input parameter of type INTEGER. The default is 0.

SRIMassDelLimit
Specifies a criterion for recommending that the RUNSTATS utility should be run on an index space. If the number of mass deletes from an index space or partition since the last REORG, REBUILD INDEX, or LOAD REPLACE is greater than this value, DSNACCOR recommends running RUNSTATS. (See item 3 in Figure 50 on page 668.) This is an input parameter of type INTEGER. The default is 0.

ExtentLimit
Specifies a criterion for recommending that the RUNSTATS or REORG utility

should be run on a table space or index space. Also specifies that DSNACCOR should warn the user that the table space or index space has used too many extents. If the number of physical extents in the index space, table space, or partition is greater than this value, DSNACCOR recommends running RUNSTATS or REORG and altering data set allocations. (See Figure 51 on page 668.) This is an input parameter of type INTEGER. The default is 50.

LastStatement
When DSNACCOR returns a severe error (return code 12), this field contains the SQL statement that was executing when the error occurred. This is an output parameter of type VARCHAR(8012).

ReturnCode
The return code from DSNACCOR execution. Possible values are:

0 DSNACCOR executed successfully. The *ErrorMsg* parameter contains the approximate percentage of the total number of objects in the subsystem that have information in the real-time statistics tables.

4 DSNACCOR completed, but one or more input parameters might be incompatible. The *ErrorMsg* parameter contains the input parameters that might be incompatible.

8 DSNACCOR terminated with errors. The *ErrorMsg* parameter contains a message that describes the error.

12 DSNACCOR terminated with severe errors. The *ErrorMsg* parameter contains a message that describes the error. The *LastStatement* parameter contains the SQL statement that was executing when the error occurred.

14 DSNACCOR terminated because it could not access one or more of the real-time statistics tables. The *ErrorMsg* parameter contains the names of the tables that DSNACCOR could not access.

15 DSNACCOR terminated because it encountered a problem with one of the declared temporary tables that it defines and uses.

16 DSNACCOR terminated because it could not define a declared temporary table. No table spaces were defined in the TEMP database.

NULL DSNACCOR terminated but could not set a return code.

This is an output parameter of type INTEGER.

ErrorMsg
Contains information about DSNACCOR execution. If DSNACCOR runs successfully (*ReturnCode*=0), this field contains the approximate percentage of objects in the subsystem that are in the real-time statistics tables. Otherwise, this field contains error messages. This is an output parameter of type VARCHAR(1331).

IFCARetCode
Contains the return code from an IFI COMMAND call. DSNACCOR issues commands through the IFI interface to determine the status of objects. This is an output parameter of type INTEGER.

IFCAResCode
Contains the reason code from an IFI COMMAND call. This is an output parameter of type INTEGER.

XSBytes
Contains the number of bytes of information that did not fit in the IFI return area after an IFI COMMAND call. This is an output parameter of type INTEGER.

Formulas for recommending actions

The following formulas specify the criteria that DSNACCOR uses for its recommendations and warnings. The variables in italics are DSNACCOR input parameters. The capitalized variables are columns of the SYSIBM.TABLESPACESTATS or SYSIBM.INDEXSPACESTATS tables. The numbers to the right of selected items are reference numbers for the option descriptions in the previous section.

```
((QueryType='COPY' OR QueryType='ALL') AND
 (ObjectType='TS' OR ObjectType='ALL') AND
 ICType='F') AND
 (COPYLASTTIME IS NULL OR
 REORGLASTTIME>COPYLASTTIME OR
 LOADRLASTTIME>COPYLASTTIME OR
 (CURRENT DATE-COPYLASTTIME)>CRDaySncLastCopy OR      1
 (COPYUPDATEDPAGES*100)/NACTIVE>CRUpdatedPagesPct OR  2
 (COPYCHANGES*100)/TOTALROWS>CRChangesPct)            3
```

Figure 44. When DSNACCOR recommends a full image copy on a table space

```
((QueryType='COPY' OR QueryType='ALL') AND
 (ObjectType='IX' OR ObjectType='ALL') AND
 (ICType='F'  OR ICType='B')) AND
 (COPYLASTTIME IS NULL OR
 REORGLASTTIME>COPYLASTTIME OR
 LOADRLASTTIME>COPYLASTTIME OR
 REBUILDLASTTIME>COPYLASTTIME OR
 (CURRENT DATE-COPYLASTTIME)>CRDaySncLastCopy OR      1
 (NACTIVE>CRIndexSize AND                             2
 ((COPYUPDATEDPAGES*100)/NACTIVE>CRUpdatedPagesPct OR 3
 (COPYCHANGES*100)/TOTALENTRIES>CRChangesPct)))       4
```

Figure 45. When DSNACCOR recommends a full image copy on an index space

```
((QueryType='COPY' OR QueryType='ALL') AND
 (ObjectType='TS' OR ObjectType='ALL') AND
 ICType='I' AND
 COPYLASTTIME IS NOT NULL) AND
 (LOADRLASTTIME>COPYLASTTIME OR
 REORGLASTTIME>COPYLASTTIME OR
 (COPYUPDATEDPAGES*100)/NACTIVE>ICRUpdatedPagesPct OR 1
 (COPYCHANGES*100)/TOTALROWS>ICRChangesPct))          2
```

Figure 46. When DSNACCOR recommends an incremental image copy on a table space

```
((QueryType='REORG' OR QueryType='ALL') AND
 (ObjectType='TS' OR ObjectType='ALL')) AND
 (REORGLASTTIME IS NULL OR
 ((REORGINSERTS+REORGDELETES+REORGUPDATES)*100)/TOTALROWS>RRTInsDelUpdPct OR  1
 (REORGUNCLUSTINS*100)/TOTALROWS>RRTUnclustInsPct OR                          2
 (REORGDISORGLOB*100)/TOTALROWS>RRTDisorgLOBPct OR                            3
 ((REORGNEARINDREF+REORGFARINDREF)*100)/TOTALROWS>RRTIndRefLimit OR           4
 REORGMASSDELETE>RRTMassDelLimit OR                                           5
 EXTENTS>ExtentLimit)                                                         6
```

Figure 47. When DSNACCOR recommends a REORG on a table space

```
((QueryType='REORG' OR QueryType='ALL') AND
 (ObjectType='IX' OR ObjectType='ALL')) AND
 (REORGLASTTIME IS NULL OR
 ((REORGINSERTS+REORGDELETES)*100)/TOTALENTRIES>RRIInsertDeletePct OR        1
 (REORGAPPENDINSERT*100)/TOTALENTRIES>RRIAppendInsertPct OR                  2
 (REORGPSEUDODELETES*100)/TOTALENTRIES>RRIPseudoDeletePct OR                 3
 REORGMASSDELETE>RRIMassDeleteLimit OR                                       4
 (REORGLEAFFAR*100)/NACTIVE>RRILeafLimit OR                                  5
 REORGNUMLEVELS>RRINumLevelsLimit OR                                         6
 EXTENTS>ExtentLimit)                                                        7
```

Figure 48. When DSNACCOR recommends a REORG on an index space

```
((QueryType='RUNSTATS' OR QueryType='ALL') AND
 (ObjectType='TS' OR ObjectType='ALL')) AND
 (STATSLASTTIME IS NULL OR
 (((STATSINSERTS+STATSDELETES+STATSUPDATES)*100)/TOTALROWS>SRTInsDelUpdPct AND  1
 (STATSINSERTS+STATSDELETES+STATSUPDATES)>SRTInsDelUpdAbs) OR                   2
 STATSMASSDELETE>SRTMassDeleteLimit)                                            3
```

Figure 49. When DSNACCOR recommends RUNSTATS on a table space

```
((QueryType='RUNSTATS' OR QueryType='ALL') AND
 (ObjectType='IX' OR ObjectType='ALL')) AND
 (STATSLASTTIME IS NULL OR
 (((STATSINSERTS+STATSDELETES)*100)/TOTALENTRIES>SRIInsDelUpdPct AND  1
 (STATSINSERTS+STATSDELETES)>SRIInsDelUpdAbs) OR                      2
 STATSMASSDELETE>SRIMassDelLimit)                                     3
```

Figure 50. When DSNACCOR recommends RUNSTATS on an index space

```
EXTENTS>ExtentLimit
```

Figure 51. When DSNACCOR warns that too many data set extents for a table space or index space are used

Using an exception table

An exception table is an optional, user-created DB2 table that you can use to place information in the INEXCEPTTABLE column of the recommendations result set. You can put any information in the INEXCEPTTABLE column, but the most common use

of this column is to filter the recommendations result set. Each row in the exception table represents an object for which you want to provide information for the recommendations result set.

To create the exception table, execute the following SQL statement:

```
CREATE TABLE DSNACC.EXCEPT_TBL
 (DBNAME CHAR(8) NOT NULL,
  NAME CHAR(8) NOT NULL,
  QUERYTYPE CHAR(40) NOT NULL);
```

The meanings of the columns are:

DBNAME
The database name for an object in the exception table.

NAME
The table space name or index space name for an object in the exception table.

QUERYTYPE
The information that you want to place in the INEXCEPTTABLE column of the recommendations result set.

If you put a null value in this column, DSNACCOR puts the value YES in the INEXCEPTTABLE column of the recommendations result set row for the object that matches the DBNAME and NAME values.

After you create the exception table, insert a row for each object for which you want to include information in the INEXCEPTTABLE column. For example, suppose that you want the INEXCEPTTABLE column to contain the string 'IRRELEVANT' for table space STAFF in database DSNDB04. You also want the INEXCEPTTABLE column to contain 'CURRENT' for table space DSN8S71D in database DSN8D71A. Execute these INSERT statements:

```
INSERT INTO DSNACC.EXCEPT_TBL VALUES('DSNDB04 ', 'STAFF ', 'IRRELEVANT');
INSERT INTO DSNACC.EXCEPT_TBL VALUES('DSN8D71A', 'DSN8S71D', 'CURRENT');
```

To use the contents of INEXCEPTTABLE for filtering, include a condition that involves the INEXCEPTTABLE column in the search condition that you specify in your *criteria* input parameter. For example, suppose that you want to include all rows for database DSNDB04 in the recommendations result set, except for those rows that contain the string 'IRRELEVANT' in the INEXCEPTTABLE column. You might include the following search condition in your *criteria* input parameter:

```
DBNAME='DSNDB04' AND INEXCEPTTABLE<>'IRRELEVANT'
```

Example of DSNACCOR invocation

The following COBOL example shows variable declarations and an SQL CALL for obtaining recommendations for objects in databases DSN8D71A and DSN8D71L.

```
 WORKING-STORAGE SECTION.
.
.
.

***********************
* DSNACCOR PARAMETERS *
***********************
 01  QUERYTYPE.
     49   QUERYTYPE-LN        PICTURE S9(4) COMP VALUE 40.
     49   QUERYTYPE-DTA       PICTURE X(40)  VALUE 'ALL'.
 01  OBJECTTYPE.
     49   OBJECTTYPE-LN       PICTURE S9(4) COMP VALUE 3.
     49   OBJECTTYPE-DTA      PICTURE X(3)  VALUE 'ALL'.
 01  ICTYPE.
```

```
#                          49   ICTYPE-LN           PICTURE S9(4) COMP VALUE 1.
#                          49   ICTYPE-DTA          PICTURE X(1)  VALUE 'B'.
#                       01  STATSSCHEMA.
#                          49   STATSSCHEMA-LN      PICTURE S9(4) COMP VALUE 128.
#                          49   STATSSCHEMA-DTA     PICTURE X(128)  VALUE 'SYSIBM'.
#                       01  CATLGSCHEMA.
#                       49   CATLGSCHEMA-LN      PICTURE S9(4) COMP VALUE 128.
#                          49   CATLGSCHEMA-DTA     PICTURE X(128)  VALUE 'SYSIBM'.
#                       01  LOCALSCHEMA.
#                          49   LOCALSCHEMA-LN      PICTURE S9(4) COMP VALUE 128.
#                          49   LOCALSCHEMA-DTA     PICTURE X(128)  VALUE 'DSNACC'.
#                       01  CHKLVL                  PICTURE S9(9) COMP VALUE +3.
#                       01  CRITERIA.
#                          49   CRITERIA-LN         PICTURE S9(4) COMP VALUE 4096.
#                          49   CRITERIA-DTA        PICTURE X(4096)  VALUE SPACES.
#                       01  RESTRICTED.
#                          49   RESTRICTED-LN       PICTURE S9(4) COMP VALUE 80.
#                          49   RESTRICTED-DTA      PICTURE X(80)  VALUE SPACES.
#                       01  CRUPDATEDPAGESPCT       PICTURE S9(9) COMP VALUE +0.
#                       01  CRCHANGESPCT            PICTURE S9(9) COMP VALUE +0.
#                       01  CRDAYSNCLASTCOPY        PICTURE S9(9) COMP VALUE +0.
#                       01  ICRUPDATEDPAGESPCT      PICTURE S9(9) COMP VALUE +0.
#                       01  ICRCHANGESPCT           PICTURE S9(9) COMP VALUE +0.
#                       01  CRINDEXSIZE             PICTURE S9(9) COMP VALUE +0.
#                       01  RRTINSDELUPDPCT         PICTURE S9(9) COMP VALUE +0.
#                       01  RRTUNCLUSTINSPCT        PICTURE S9(9) COMP VALUE +0.
#                       01  RRTDISORGLOBPCT         PICTURE S9(9) COMP VALUE +0.
#                       01  RRTMASSDELLIMIT         PICTURE S9(9) COMP VALUE +0.
#                       01  RRTINDREFLIMIT          PICTURE S9(9) COMP VALUE +0.
#                       01  RRIINSERTDELETEPCT      PICTURE S9(9) COMP VALUE +0.
#                       01  RRIAPPENDINSERTPCT      PICTURE S9(9) COMP VALUE +0.
#                       01  RRIPSEUDODELETEPCT      PICTURE S9(9) COMP VALUE +0.
#                       01  RRIMASSDELLIMIT         PICTURE S9(9) COMP VALUE +0.
#                       01  RRILEAFLIMIT            PICTURE S9(9) COMP VALUE +0.
#                       01  RRINUMLEVELSLIMIT       PICTURE S9(9) COMP VALUE +0.
#                       01  SRTINSDELUPDPCT         PICTURE S9(9) COMP VALUE +0.
#                       01  SRTINSDELUPDABS         PICTURE S9(9) COMP VALUE +0.
#                       01  SRTMASSDELLIMIT         PICTURE S9(9) COMP VALUE +0.
#                       01  SRIINSDELUPDPCT         PICTURE S9(9) COMP VALUE +0.
#                       01  SRIINSDELUPDABS         PICTURE S9(9) COMP VALUE +0.
#                       01  SRIMASSDELLIMIT         PICTURE S9(9) COMP VALUE +0.
#                       01  EXTENTLIMIT             PICTURE S9(9) COMP VALUE +0.
#                       01  LASTSTATEMENT.
#                          49   LASTSTATEMENT-LN    PICTURE S9(4) COMP VALUE 8012.
#                          49   LASTSTATEMENT-DTA   PICTURE X(8012)  VALUE SPACES.
#                       01  RETURNCODE              PICTURE S9(9) COMP VALUE +0.
#                       01  ERRORMSG
#                          49   ERRORMSG-LN         PICTURE S9(4) COMP VALUE 1331.
#                          49   ERRORMSG-DTA        PICTURE X(1331)  VALUE SPACES.
#                       01  IFCARETCODE             PICTURE S9(9) COMP VALUE +0.
#                       01  IFCARESCODE             PICTURE S9(9) COMP VALUE +0.
#                       01  XSBYTES                 PICTURE S9(9) COMP VALUE +0.
#                      ****************************************
#                      * INDICATOR VARIABLES.                 *
#                      * INITIALIZE ALL NON-ESSENTIAL INPUT    *
#                      * VARIABLES TO -1, TO INDICATE THAT THE *
#                      * INPUT VALUE IS NULL.                 *
#                      ****************************************
#                       01  QUERYTYPE-IND           PICTURE S9(4) COMP-4 VALUE +0.
#                       01  OBJECTTYPE-IND          PICTURE S9(4) COMP-4 VALUE +0.
#                       01  ICTYPE-IND              PICTURE S9(4) COMP-4 VALUE +0.
#                       01  STATSSCHEMA-IND         PICTURE S9(4) COMP-4 VALUE -1.
#                       01  CATLGSCHEMA-IND         PICTURE S9(4) COMP-4 VALUE -1.
#                       01  LOCALSCHEMA-IND         PICTURE S9(4) COMP-4 VALUE -1.
#                       01  CHKLVL-IND              PICTURE S9(4) COMP-4 VALUE -1.
#                       01  CRITERIA-IND            PICTURE S9(4) COMP-4 VALUE -1.
#                       01  RESTRICTED-IND          PICTURE S9(4) COMP-4 VALUE -1.
```

```
#                      01  CRUPDATEDPAGESPCT-IND    PICTURE S9(4) COMP-4 VALUE -1.
#                      01  CRCHANGESPCT-IND         PICTURE S9(4) COMP-4 VALUE -1.
#                      01  CRDAYSNCLASTCOPY-IND     PICTURE S9(4) COMP-4 VALUE -1.
#                      01  ICRUPDATEDPAGESPCT-IND   PICTURE S9(4) COMP-4 VALUE -1.
#                      01  ICRCHANGESPCT-IND        PICTURE S9(4) COMP-4 VALUE -1.
#                      01  CRINDEXSIZE-IND          PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRTINSDELUPDPCT-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRTUNCLUSTINSPCT-IND     PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRTDISORGLOBPCT-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRTMASSDELLIMIT-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRTINDREFLIMIT-IND       PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRIINSERTDELETEPCT-IND   PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRIAPPENDINSERTPCT-IND   PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRIPSEUDODELETEPCT-IND   PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRIMASSDELLIMIT-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRILEAFLIMIT-IND         PICTURE S9(4) COMP-4 VALUE -1.
#                      01  RRINUMLEVELSLIMIT-IND    PICTURE S9(4) COMP-4 VALUE -1.
#                      01  SRTINSDELUPDPCT-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  SRTINSDELUPDABS-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  SRTMASSDELLIMIT-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  SRIINSDELUPDPCT-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  SRIINSDELUPDABS-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  SRIMASSDELLIMIT-IND      PICTURE S9(4) COMP-4 VALUE -1.
#                      01  EXTENTLIMIT-IND          PICTURE S9(4) COMP-4 VALUE -1.
#                      01  LASTSTATEMENT-IND        PICTURE S9(4) COMP-4 VALUE +0.
#                      01  RETURNCODE-IND           PICTURE S9(4) COMP-4 VALUE +0.
#                      01  ERRORMSG-IND             PICTURE S9(4) COMP-4 VALUE +0.
#                      01  IFCARETCODE-IND          PICTURE S9(4) COMP-4 VALUE +0.
#                      01  IFCARESCODE-IND          PICTURE S9(4) COMP-4 VALUE +0.
#                      01  XSBYTES-IND              PICTURE S9(4) COMP-4 VALUE +0.
#
#                      PROCEDURE DIVISION.
#                     :
#                     :
#
#                     ***********************************************************
#                     * SET VALUES FOR DSNACCOR INPUT PARAMETERS:               *
#                     * - USE THE CHKLVL PARAMETER TO CAUSE DSNACCOR TO CHECK *
#                     *   FOR ORPHANED OBJECTS AND INDEX SPACES WITHOUT       *
#                     *   TABLE SPACES, BUT INCLUDE THOSE OBJECTS IN THE      *
#                     *   RECOMMENDATIONS RESULT SET (CHKLVL=1+2+16=19)       *
#                     * - USE THE CRITERIA PARAMETER TO CAUSE DSNACCOR TO     *
#                     *   MAKE RECOMMENDATIONS ONLY FOR OBJECTS IN DATABASES  *
#                     *   DSN8D71A AND DSN8D71L.                              *
#                     * - FOR THE FOLLOWING PARAMETERS, SET THESE VALUES,     *
#                     *   WHICH ARE LOWER THAN THE DEFAULTS:                  *
#                     *   CRUPDATEDPAGESPCT   4                               *
#                     *   CRCHANGESPCT        2                               *
#                     *   RRTINSDELUPDPCT     2                               *
#                     *   RRTUNCLUSTINSPCT    5                               *
#                     *   RRTDISORGLOBPCT     5                               *
#                     *   RRIAPPENDINSERTPCT  5                               *
#                     *   SRTINSDELUPDPCT     5                               *
#                     *   SRIINSDELUPDPCT     5                               *
#                     *   EXTENTLIMIT         3                               *
#                     ***********************************************************
#                          MOVE 19 TO CHKLVL.
#                          MOVE SPACES TO CRITERIA-DTA.
#                          MOVE 'DBNAME = ''DSN8D71A'' OR DBNAME = ''DSN8D71L'''
#                                TO CRITERIA-DTA.
#                          MOVE 46 TO CRITERIA-LN.
#                          MOVE 4 TO CRUPDATEDPAGESPCT.
#                          MOVE 2 TO CRCHANGESPCT.
#                          MOVE 2 TO RRTINSDELUPDPCT.
#                          MOVE 5 TO RRTUNCLUSTINSPCT.
#                          MOVE 5 TO RRTDISORGLOBPCT.
#                          MOVE 5 TO RRIAPPENDINSERTPCT.
#                          MOVE 5 TO SRTINSDELUPDPCT.
```

```
#                          MOVE 5 TO SRIINSDELUPDPCT.
#                          MOVE 3 TO EXTENTLIMIT.
#                        *******************************
#                        * INITIALIZE OUTPUT PARAMETERS *
#                        *******************************
#                          MOVE SPACES TO LASTSTATEMENT-DTA.
#                          MOVE 1 TO LASTSTATEMENT-LN.
#                          MOVE 0 TO RETURNCODE-O2.
#                          MOVE SPACES TO ERRORMSG-DTA.
#                          MOVE 1 TO ERRORMSG-LN.
#                          MOVE 0 TO IFCARETCODE.
#                          MOVE 0 TO IFCARESCODE.
#                          MOVE 0 TO XSBYTES.
#                        ******************************************************
#                        * SET THE INDICATOR VARIABLES TO 0 FOR NON-NULL INPUT *
#                        * PARAMETERS (PARAMETERS FOR WHICH YOU DO NOT WANT    *
#                        * DSNACCOR TO USE DEFAULT VALUES) AND FOR OUTPUT      *
#                        * PARAMETERS.                                        *
#                        ******************************************************
#                          MOVE 0 TO CHKLVL-IND.
#                          MOVE 0 TO CRITERIA-IND.
#                          MOVE 0 TO CRUPDATEDPAGESPCT-IND.
#                          MOVE 0 TO CRCHANGESPCT-IND.
#                          MOVE 0 TO RRTINSDELUPDPCT-IND.
#                          MOVE 0 TO RRTUNCLUSTINSPCT-IND.
#                          MOVE 0 TO RRTDISORGLOBPCT-IND.
#                          MOVE 0 TO RRIAPPENDINSERTPCT-IND.
#                          MOVE 0 TO SRTINSDELUPDPCT-IND.
#                          MOVE 0 TO SRIINSDELUPDPCT-IND.
#                          MOVE 0 TO EXTENTLIMIT-IND.
#                          MOVE 0 TO LASTSTATEMENT-IND.
#                          MOVE 0 TO RETURNCODE-IND.
#                          MOVE 0 TO ERRORMSG-IND.
#                          MOVE 0 TO IFCARETCODE-IND.
#                          MOVE 0 TO IFCARESCODE-IND.
#                          MOVE 0 TO XSBYTES-IND.
#                        .
#
#
#                        *****************
#                        * CALL DSNACCOR *
#                        *****************
#                          EXEC SQL
#                           CALL SYSPROC.DSNACCOR
#                           (:QUERYTYPE           :QUERYTYPE-IND,
#                            :OBJECTTYPE          :OBJECTTYPE-IND,
#                            :ICTYPE              :ICTYPE-IND,
#                            :STATSSCHEMA         :STATSSCHEMA-IND,
#                            :CATLGSCHEMA         :CATLGSCHEMA-IND,
#                            :LOCALSCHEMA         :LOCALSCHEMA-IND,
#                            :CHKLVL              :CHKLVL-IND,
#                            :CRITERIA            :CRITERIA-IND,
#                            :CRUPDATEDPAGESPCT   :CRUPDATEDPAGESPCT-IND,
#                            :CRCHANGESPCT        :CRCHANGESPCT-IND,
#                            :CRDAYSNCLASTCOPY    :CRDAYSNCLASTCOPY-IND,
#                            :ICRUPDATEDPAGESPCT  :ICRUPDATEDPAGESPCT-IND,
#                            :CRINDEXSIZE         :CRINDEXSIZE-IND,
#                            :RRTINSDELUPDPCT     :RRTINSDELUPDPCT-IND,
#                            :RRTUNCLUSTINSPCT    :RRTUNCLUSTINSPCT-IND,
#                            :RRTDISORGLOBPCT     :RRTDISORGLOBPCT-IND,
#                            :RRTMASSDELLIMIT     :RRTMASSDELLIMIT-IND,
#                            :RRTINDREFLIMIT      :RRTINDREFLIMIT-IND,
#                            :RRIINSERTDELETEPCT  :RRIINSERTDELETEPCT-IND,
#                            :RRIAPPENDINSERTPCT  :RRIAPPENDINSERTPCT-IND,
#                            :RRIPSEUDODELETEPCT  :RRIPSEUDODELETEPCT-IND,
#                            :RRIMASSDELLIMIT     :RRIMASSDELLIMIT-IND,
#                            :RRILEAFLIMIT        :RRILEAFLIMIT-IND,
```

```
                    :RRINUMLEVELSLIMIT   :RRINUMLEVELSLIMIT-IND,
                    :SRTINSDELUPDPCT     :SRTINSDELUPDPCT-IND,
                    :SRTINSDELUPDABS     :SRTINSDELUPDABS-IND,
                    :SRTMASSDELLIMIT     :SRTMASSDELLIMIT-IND,
                    :SRIINSDELUPDPCT     :SRIINSDELUPDPCT-IND,
                    :SRIINSDELUPDABS     :SRIINSDELUPDABS-IND,
                    :SRIMASSDELLIMIT     :SRIMASSDELLIMIT-IND,
                    :EXTENTLIMIT         :EXTENTLIMIT-IND,
                    :LASTSTATEMENT       :LASTSTATEMENT-IND,
                    :RETURNCODE          :RETURNCODE-IND,
                    :ERRORMSG            :ERRORMSG-IND,
                    :IFCARETCODE         :IFCARETCODE-IND,
                    :IFCARESCODE         :IFCARESCODE-IND,
                    :XSBYTES             :XSBYTES-IND)
                  END-EXEC.
```

DSNACCOR output

If DSNACCOR executes successfully, in addition to the output parameters described in "DSNACCOR option descriptions" on page 660, DSNACCOR returns two result sets.

The first result set contains the results from IFI COMMAND calls that DSNACCOR makes. Table 125 shows the format of the first result set.

Table 125. Result set row for first DSNACCOR result set

Column name	Data type	Contents
RS_SEQUENCE	INTEGER	Sequence number of the output line
RS_DATA	CHAR(80)	A line of command output

The second result set contains DSNACCOR's recommendations. This result set contains one or more rows for a table space or index space. A nonpartitioned table space or nonpartitioning index space can have at most one row in the result set. A partitioned table space or partitioning index space can have at most one row for each partition. A table space, index space, or partition has a row in the result set if the following conditions are true:

- If the *Criteria* input parameter contains a search condition, the search condition is true for the table space, index space, or partition.
- DSNACCOR recommends at least one action for the table space, index space, or partition.

Table 126 shows the columns of a result set row.

Table 126. Result set row for second DSNACCOR result set

Column name	Data type	Description
DBNAME	CHAR(8)	Name of the database that contains the object.
NAME	CHAR(8)	Table space or index space name.
PARTITION	INTEGER	Data set number or partition number.
OBJECTTYPE	CHAR(2)	DB2 object type: • TS for a table space • IX for an index space

Table 126. Result set row for second DSNACCOR result set (continued)

Column name	Data type	Description
OBJECTSTATUS	CHAR(36)	Status of the object: • ORPHANED, if the object is an index space with no corresponding table space, or the object does not exist • If the object is in a restricted state, one of the following values: – TS=*restricted-state*, if OBJECTTYPE is TS – IX=*restricted-state*, if OBJECTTYPE is IX *restricted-state* is one of the status codes that appear in DISPLAY DATABASE output. See Chapter 2 of *DB2 Command Reference* for details.
IMAGECOPY	CHAR(3)	COPY recommendation: • If OBJECTTYPE is TS: FUL (full image copy), INC (incremental image copy), or NO • If OBJECTTYPE is IX: YES or NO
RUNSTATS	CHAR(3)	RUNSTATS recommendation: YES or NO.
EXTENTS	CHAR(3)	Whether the data sets for the object have exceeded *ExtentLimit*: YES or NO.
REORG	CHAR(3)	REORG recommendation: YES or NO.
INEXCEPTTABLE	CHAR(40)	A string that contains one of the following values: • Text that you specify in the QUERYTYPE column of the exception table. • YES, if you put a row in the exception table for the object that this result set row represents, but you specify NULL in the QUERYTYPE column. • NO, if the exception table exists but does not have a row for the object that this result set row represents. • Null, if the exception table does not exist, or the *ChkLvl* input parameter does not include the value 4.
ASSOCIATEDTS	CHAR(8)	If OBJECTTYPE is IX and the *ChkLvl* input parameter includes the value 2, this value is the name of the table space that is associated with the index space. Otherwise this value is null.
COPYLASTTIME	TIMESTAMP	Timestamp of the last full image copy on the object. This value is null if COPY was never run, or the last COPY execution was terminated.
LOADRLASTTIME	TIMESTAMP	Timestamp of the last LOAD REPLACE on the object. NULL if LOAD REPLACE was never run, or the last LOAD REPLACE execution was terminated.
REBUILDLASTTIME	TIMESTAMP	Timestamp of the last REBUILD INDEX on the object. This value is null if REBUILD INDEX was never run, or if the last REBUILD INDEX execution was terminated.
CRUPDPGSPCT	INTEGER	If OBJECTTYPE is TS or IX and IMAGECOPY is YES, the ratio of distinct updated pages to preformatted pages, expressed as a percentage. Otherwise, this value is null.
CRCPYCHGPCT	INTEGER	If OBJECTTYPE is TS and IMAGECOPY is YES, this value is the ratio of the number INSERTs, UPDATEs, and DELETEs since the last image copy to the total number of rows or LOBs in the table space or partition, expressed as a percentage. If OBJECTTYPE is IX and IMAGECOPY is YES, this value is the ratio of the number INSERTs and DELETEs since the last image copy to the total number of entries in the index space or partition, expressed as a percentage. Otherwise, this value is null.

Table 126. Result set row for second DSNACCOR result set (continued)

Column name	Data type	Description
CRDAYSCELSTCPY	INTEGER	If OBJECTTYPE is TS or IX and IMAGECOPY is YES, the number of days since the last image copy. Otherwise, this value is null.
CRINDEXSIZE	INTEGER	If OBJECTTYPE is IX and IMAGECOPY is YES, the number of active pages in the index space or partition. Otherwise, this value is null.
REORGLASTTIME	TIMESTAMP	Timestamp of the last REORG on the object. This value is null if REORG was never run, or if the last REORG execution was terminated.
RRTINSDELUPDPCT	INTEGER	If OBJECTTYPE is TS and REORG is YES, the ratio of the sum of INSERTs, UPDATEs, and DELETEs since the last REORG to the total number of rows or LOBs in the table space or partition, expressed as a percentage. Otherwise, this value is null.
RRTUNCINSPCT	INTEGER	If OBJECTTYPE is TS and REORG is YES, the ratio of the number of unclustered INSERTs to the total number of rows or LOBs in the table space or partition, expressed as a percentage. Otherwise, this value is null.
RRTDISORGLOBPCT	INTEGER	If OBJECTTYPE is TS and REORG is YES, the ratio of the number of imperfectly chunked LOBs to the total number of rows or LOBs in the table space or partition, expressed as a percentage. Otherwise, this value is null.
RRTMASSDELETE	INTEGER	If OBJECTTYPE is TS, REORG is YES, and the table space is a segmented table space or LOB table space, this value is the number of mass deletes since the last REORG or LOAD REPLACE. If OBJECTTYPE is TS, REORG is YES, and the table space is nonsegmented, this value is the number of dropped tables since the last REORG or LOAD REPLACE. Otherwise, this value is null.
RRTINDREF	INTEGER	If OBJECTTYPE is TS, REORG is YES, the ratio of the total number of overflow records that were created since the last REORG or LOAD REPLACE to the total number of rows or LOBs in the table space or partition, expressed as a percentage. Otherwise, this value is null.
RRIINSDELPCT	INTEGER	If OBJECTTYPE is IX and REORG is YES, the ratio of the sum of INSERTs and DELETEs since the last REORG to the total number of index entries in the index space or partition, expressed as a percentage. Otherwise, this value is null.
RRIAPPINSPCT	INTEGER	If OBJECTTYPE is IX and REORG is YES, the ratio of the number of index entries that were inserted since the last REORG, REBUILD INDEX, or LOAD REPLACE and had a key value greater than the maximum key value in the index space or partition, to the number of index entries in the index space or partition, expressed as a percentage. Otherwise, this value is null.
RRIPSDDELPCT	INTEGER	If OBJECTTYPE is IX and REORG is YES, the ratio of the number of index entries that were pseudo-deleted since the last REORG, REBUILD INDEX, or LOAD REPLACE to the number of index entries in the index space or partition, expressed as a percentage. Otherwise, this value is null.

Table 126. Result set row for second DSNACCOR result set (continued)

Column name	Data type	Description
RRIMASSDELETE	INTEGER	If OBJECTTYPE is IX and REORG is YES, the number of mass deletes from the index space or partition since the last REORG, REBUILD, or LOAD REPLACE. Otherwise, this value is null.
RRILEAF	INTEGER	If OBJECTTYPE is IX and REORG is YES, the ratio of the number of index page splits that occurred since the last REORG, REBUILD INDEX, or LOAD REPLACE in which the higher part of the split page was far from the location of the original page, to the total number of active pages in the index space or partition, expressed as a percentage. Otherwise, this value is null.
RRINUMLEVELS	INTEGER	If OBJECTTYPE is IX and REORG is YES, the number of levels in the index tree that were added or removed since the last REORG, REBUILD INDEX, or LOAD REPLACE. Otherwise, this value is null.
STATSLASTTIME	TIMESTAMP	Timestamp of the last RUNSTATS on the object. This value is null if RUNSTATS was never run, or if the last RUNSTATS execution was terminated.
SRTINSDELPCT	INTEGER	If OBJECTTYPE is TS and RUNSTATS is YES, the ratio of the number INSERTs, UPDATEs, and DELETEs since the last RUNSTATS on a table space or partition, to the total number of rows or LOBs in the table space or partition, expressed as a percentage. Otherwise, this value is null.
SRTINSDELABS	INTEGER	If OBJECTTYPE is TS and RUNSTATS is YES, the number INSERTs, UPDATEs, and DELETEs since the last RUNSTATS on a table space or partition. Otherwise, this value is null.
SRTMASSDELETE	INTEGER	If OBJECTTYPE is TS and RUNSTATS is YES, the number of mass deletes from the table space or partition since the last REORG or LOAD REPLACE. Otherwise, this value is null.
SRIINSDELPCT	INTEGER	If OBJECTTYPE is IX and RUNSTATS is YES, the ratio of the number INSERTs and DELETEs since the last RUNSTATS on the index space or partition, to the total number of index entries in the index space or partition, expressed as a percentage. Otherwise, this value is null.
SRIINSDELABS	INTEGER	If OBJECTTYPE is IX and RUNSTATS is YES, the number INSERTs and DELETEs since the last RUNSTATS on the index space or partition. Otherwise, this value is null.
SRIMASSDELETE	INTEGER	If OBJECTTYPE is IX and RUNSTATS is YES, the number of mass deletes from the index space or partition since the last REORG, REBUILD INDEX, or LOAD REPLACE. Otherwise, this value is null.
TOTALEXTENTS	SMALLINT	If EXTENTS is YES, the number of physical extents in the table space, index space, or partition. Otherwise, this value is null.

Appendix C. Resetting an advisory or restrictive status

DB2 sets a restrictive or advisory status on an object to control access and help ensure data integrity. This appendix outlines the restrictive and non-restrictive object states that affect utilities, and the steps required to correct each status for a particular object.

Use the DISPLAY DATABASE command to display the current status for an object.

The following states are described in this section:

"Auxiliary CHECK-pending status"
"Auxiliary warning status"
"CHECK-pending status" on page 678
"COPY-pending status" on page 679
"Group buffer pool RECOVER-pending status" on page 679
"Informational COPY-pending status" on page 680
"REBUILD-pending status" on page 680
"RECOVER-pending status" on page 681
"REFRESH-pending status" on page 682
"REORG-pending status" on page 682
"Restart pending status" on page 683

Auxiliary CHECK-pending status

The auxiliary CHECK-pending restrictive status is set on when at least one base table LOB column error is detected and not invalidated as a result of running CHECK DATA AUXERROR REPORT.

Refer to Table 127 for information about resetting the auxiliary CHECK-pending status.

Table 127. Resetting auxiliary CHECK-pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
Auxiliary CHECK-pending	ACHKP	base table space	1. Update or delete invalid LOBs using SQL. 2. Run the CHECK DATA utility with the appropriate SCOPE option to verify the validity of LOBs and reset ACHKP status. You can use the REPAIR utility followed by CHECK DATA to reset the ACHKP status, but use caution.	1

Notes:

1. A base table space in the ACHKP status is unavailable for processing by SQL.

Auxiliary warning status

Auxiliary warning (AUXW) status is set on when at least one base table LOB column has an invalidated LOB as a result of running CHECK DATA AUXERROR INVALIDATE. An attempt to retrieve an invalidated LOB results in a -904 SQL return code.

The RECOVER utility also sets AUXW status if it finds an invalid LOB column. Invalid LOB columns might result from the following actions (all three must apply):

© Copyright IBM Corp. 1983, 2001

Resetting an advisory or restrictive status

1. LOB table space was defined with LOG NO.
2. LOB table space was recovered.
3. LOB was updated since the last image copy.

Refer to Table 128 for information about resetting the auxiliary warning status.

Table 128. Resetting auxiliary warning status

Status	Abbreviation	Object Affected	Corrective Action	Notes
Auxiliary warning	AUXW	Base table space	1. Update or delete invalid LOBs using SQL. 2. Run CHECK DATA utility to verify the validity of LOBs and reset AUXW status.	1,2,3
		LOB table space	1. Update or delete invalid LOBs using SQL. 2. Run CHECK LOB utility to verify the validity of LOBs and reset AUXW status. Alternatively, you can use the REPAIR utility to set NOLOBCHKP. Be aware that using the REPAIR utility to delete invalid LOBs might cause the base table and the index on the auxiliary table to reference invalid LOBs.	1

Notes:

1. A base table space or LOB table space in the AUXW status is available for processing by SQL, even though it contains invalid LOBs. However, an attempt to retrieve an invalid LOB results in a -904 SQL return code.
2. DB2 can access all rows of a base table space that are in the AUXW status. SQL can update the invalid LOB column and delete base table rows, but the value of the LOB column cannot be retrieved. If DB2 attempts to access an invalid LOB column, a -904 SQL code is returned. The AUXW status remains on the base table space even when SQL deletes or updates the last invalid LOB column.
3. If CHECK DATA AUXERROR REPORT encounters only invalid LOB columns and no other LOB column errors, the
base table space is set to auxiliary warning status.

CHECK-pending status

The CHECK-pending restrictive status indicates that an object might be in an inconsistent state and must be checked.

The following utilities set the CHECK-pending status on a table space if referential integrity constraints are encountered:

- LOAD with ENFORCE NO
- RECOVER to a point-in-time
- CHECK LOB

The CHECK-pending status can also affect a base table space or a LOB table space.

Refer to Table 129 on page 679 for information about resetting the CHECK-pending status.

Table 129. Resetting CHECK-pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
CHECK-pending	CHKP	Table space, base table space	Check and correct referential integrity constraints using the CHECK DATA utility. If a table space is in both REORG-pending and CHECK pending status (or auxiliary CHECK-pending status), run REORG first and then CHECK DATA to clear the respective states.	
		Partitioning index, nonpartitioning index, index on the auxiliary table	1. Run CHECK INDEX on the index. 2. If any errors are found, use the REBUILD INDEX utility to rebuild the index from existing data.	1
		LOB table space	Use the CHECK LOB utility to check the LOB table space. If any errors are found: 1. Correct any defects found in the LOB table space using the REPAIR utility. 2. Run CHECK LOB again to reset the CHECK-pending status. 3. See Table 128 on page 678 if an AUXW status exists.	

Notes:

1. An index might be placed in the CHECK-pending status if you recovered an index to a specific RBA or LRSN from a copy and applied the log records, but you did not recover the table space in the same list. The CHECK-pending status can also be placed on an index if you specified the table space and the index in the same list, but the RECOVER point-in-time was not a QUIESCE or COPY SHRLEVEL REFERENCE point.

COPY-pending status

The COPY-pending restrictive status indicates that the affected object must be copied.

Refer to Table 130 for information about resetting the COPY-pending status.

Table 130. Resetting COPY-pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
COPY-pending	COPY	table space, table space partition	Take an image copy of the affected object.	

Group buffer pool RECOVER-pending status

The group buffer pool RECOVER-pending status is set on when a coupling facility crashes with pages that were not externalized. The affected object must be recovered.

Refer to Table 131 on page 680 for information about resetting the group buffer pool RECOVER-pending status.

Resetting an advisory or restrictive status

Table 131. Resetting group buffer pool RECOVER-pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
Group buffer pool RECOVER-pending	GRECP	object	Recover the object, or use START DATABASE to recover the object.	

Informational COPY-pending status

The informational COPY-pending advisory status indicates that the affected object should be copied.

Refer to Table 132 for information about resetting the informational COPY-pending status.

Table 132. Resetting informational COPY-pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
Informational COPY-pending	ICOPY	Partitioning index, nonpartitioning index, index on the auxiliary table	Copy the affected index.	

REBUILD-pending status

The REBUILD-pending (RBDP) restrictive status indicates that the affected index or index partition is broken and must be rebuilt from the data.

Refer to Table 133 for information about resetting the REBUILD-pending status.

Table 133. Resetting REBUILD-pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
REBUILD-pending	RBDP	physical or logical index partition	Run the REBUILD utility on the affected index partition.	
	RBDP*	logical partitions of nonpartitioning indexes	Run REBUILD INDEX PART or RECOVER utility on the affected logical partitions.	1
	PSRBD	nonpartitioning index, index on the auxiliary table	Run REBUILD INDEX ALL, the RECOVER utility, or run REBUILD INDEX listing all indexes in the affected index space.	

Table 133. Resetting REBUILD-pending status (continued)

Status	Abbreviation	Object Affected	Corrective Action	Notes
			The following actions also reset the REBUILD-pending status: • LOAD REPLACE the table space or partition. • REPAIR SET INDEX with NORBDPEND on the index partition. Be aware that this does not correct the data inconsistency in the index partition. Use CHECK INDEX instead of REPAIR to verify referential integrity constraints. • Start the database containing the index space with ACCESS FORCE. Be aware that this does not correct the data inconsistency in the index partition. • Run REORG INDEX SORTDATA on the affected index.	

Notes:

1. The entire nonpartitioning index is inaccessible until you reset the RBDP* status.

RECOVER-pending status

The RECOVER-pending restrictive status indicates that a table space or table space partition is broken and must be recovered.

Refer to Table 134 for information about resetting the RECOVER-pending status.

Table 134. Resetting RECOVER-pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
RECOVER-pending	RECP	table space	Run the RECOVER utility on the affected object.	
		Table space partition	Recover the partition.	
		Index on the auxiliary table	Correct the RECOVER-pending status using one of the following utilities: • REBUILD INDEX • RECOVER INDEX • REORG INDEX SORTDATA	
		Index space	In a coexistence situation, you might have an index space set in RECP status by a Version 5 subsystem that is visible on a Version 6 subsystem. Likewise, you might have an index space set in RBDP, RBDP*, or PSRBDP status by a Version 6 subsystem that is visible on a Version 5 subsystem. In either situation, run one of the following utilities on the affected index space to reset the pending status: • REBUILD INDEX • RECOVER INDEX • REORG INDEX SORTDATA	

Table 134. Resetting RECOVER-pending status (continued)

Status	Abbreviation	Object Affected	Corrective Action	Notes
			The following actions also reset the RECOVER-pending status: • LOAD REPLACE the table space or partition. • REPAIR SET TABLESPACE or INDEX with NORCVRPEND on table space or partition. Be aware that this does not correct the data inconsistency in the table space or partition. • Start the database containing the table space or index space with ACCESS FORCE. Be aware that this does not correct the data inconsistency in the table space or partition.	

REFRESH-pending status

Whenever DB2 marks an object in refresh-pending (REFP) status, it also puts the object in RECOVER-pending (RECP) or REBUILD-pending (RBDP or PSRBD). If a user-defined table space is in refresh-pending (REFP) status, you can replace the data using LOAD REPLACE. At the successful completion of the RECOVER and LOAD REPLACE job, both (REFP and RECP or REFP and RBDP or PSRBD) states will be reset.

REORG-pending status

The REORG-pending restrictive status indicates that a table space partition is broken and must be reorganized.

Refer to Table 135 for information about resetting the REORG-pending status.

Table 135. Resetting REORG-pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
REORG-pending	REORP	table space	Perform one of the following actions: • LOAD REPLACE the entire table space. • Run the REORG TABLESPACE utility with SHRLEVEL NONE. If a table space is in both REORG-pending and CHECK-pending status (or auxiliary CHECK-pending status), run REORG first and then CHECK DATA to clear the respective states. • Run REORG PART*m*:*n* SHRLEVEL NONE.	1,2, 3,4
REORG-pending	REORP	partitioned table space	For row lengths <= 32KB: 1. Run REORG TABLESPACE SHRLEVEL NONE SORTDATA. For row lengths > 32KB: 1. Run REORG TABLESPACE UNLOAD ONLY. 2. Run LOAD TABLESPACE FORMAT UNLOAD.	

Table 135. Resetting REORG-pending status (continued)

Status	Abbreviation	Object Affected	Corrective Action	Notes

Notes:

1. Consider running COPY after resetting the REORP status. Be aware that you can only use an image copy that was created before turning off the REORP status if you are performing a point-in-time recovery.
2. You cannot run SELECT, INSERT, DELETE, or UPDATE on data in a table space that is in REORP status; this includes access through a partitioning or nonpartitioning index. The only SQL access that is allowed is DROP TABLESPACE.
3. The START DATABASE ACCESS FORCE command does *not* remove the REORP status from an object.
4. You must allocate a discard data set (SYSDISC) or specify DISCARDDN if the last partition is in REORP.

Restart pending status

The restart pending status is set on if an object has back-out work pending at the end of DB2 restart.

Refer to Table 136 for information about resetting the restart pending status.

Table 136. Resetting restart pending status

Status	Abbreviation	Object Affected	Corrective Action	Notes
Restart pending	RESTP	table space, table space partitions, index spaces, and physical index space partitions	Objects in the RESTP status remain unavailable until back-out work is complete, or until restart is canceled and a conditional restart or cold start is performed in its place. See Part 4 (Volume 1) of *DB2 Administration Guide* for information about the RESTP restrictive status.	1,2,3

Notes:

1. Delay running REORG TABLESPACE SHRLEVEL CHANGE until all RESTP statuses are reset.
2. You cannot use LOAD REPLACE on an object that is in the RESTP status.
3. Utility activity against RESTP page sets or partitions is not allowed. Any attempt to access a RESTP page set or partition terminates with return code 8.

Resetting an advisory or restrictive status

Appendix D. How to run sample programs DSNTIAUL, DSNTIAD, and DSNTEP2

DB2 provides three sample programs that many users find helpful as productivity aids. These programs are shipped as source code, so you can modify them to meet your needs. The programs are:

DSNTIAUL The sample unload program. This program, which is written in assembler language, unloads some or all rows from up to 100 DB2 tables. With DSNTIAUL, you can unload data of any DB2 built-in data type or distinct type. You can unload up to 32KB of data from a LOB column. DSNTIAUL unloads the rows in a form that is compatible with the LOAD utility and generates utility control statements for LOAD. DSNTIAUL also lets you execute any SQL non-SELECT statement that can be executed dynamically.

DSNTIAD A sample dynamic SQL program in assembler language. With this program, you can execute any SQL statement that can be executed dynamically, except a SELECT statement.

DSNTEP2 A sample dynamic SQL program in the PL/I language. With this program, you can execute any SQL statement that can be executed dynamically. You can use the source version of DSNTEP2 and modify it to meet your needs, or, if you do not have a PL/I compiler at your installation, you can use the object code version of DSNTEP2.

Because these three programs also accept the static SQL statements CONNECT, SET CONNECTION, and RELEASE, you can use the programs to access DB2 tables at remote locations.

DSNTIAUL and DSNTIAD are shipped only as source code, so you must precompile, assemble, link, and bind them before you can use them. If you want to use the source code version of DSNTEP2, you must precompile, compile, link and bind it. You need to bind the object code version of DSNTEP2 before you can use it. Usually, your system administrator prepares the programs as part of the installation process. Table 137 indicates which installation job prepares each sample program. All installation jobs are in data set DSN710.SDSNSAMP.

Table 137. Jobs that prepare DSNTIAUL, DSNTIAD, and DSNTEP2

Program name	Program preparation job
DSNTIAUL	DSNTEJ2A
DSNTIAD	DSNTIJTM
DSNTEP2 (source)	DSNTEJ1P
DSNTEP2 (object)	DSNTEJ1L

To run the sample programs, use the DSN RUN command, which is described in detail in Chapter 2 of *DB2 Command Reference*. Table 138 on page 686 lists the load module name and plan name you must specify, and the parameters you can specify when you run each program. See the following sections for the meaning of each parameter.

© Copyright IBM Corp. 1983, 2001

Table 138. DSN RUN option values for DSNTIAUL, DSNTIAD, and DSNTEP2

Program name	Load module	Plan	Parameters
DSNTIAUL	DSNTIAUL	DSNTIB71	SQL
DSNTIAD	DSNTIAD	DSNTIA71	RC0 SQLTERM(*termchar*)
DSNTEP2	DSNTEP2	DSNTEP71	ALIGN(MID) or ALIGN(LHS) NOMIXED or MIXED SQLTERM(*termchar*)

The remainder of this appendix contains the following information about running each program:

- Descriptions of the input parameters
- Data sets you must allocate before you run the program
- Return codes from the program
- Examples of invocation

See the sample jobs listed in Table 137 on page 685 for a working example of each program.

Running DSNTIAUL

This section contains information that you need when you run DSNTIAUL, including parameters, data sets, return codes, and invocation examples.

DSNTIAUL parameters: DSNTIAUL accepts one parameter, SQL. If you specify this parameter, your input data set contains one or more complete SQL statements, each of which ends with a semi-colon. You can include any SQL statement that can be executed dynamically in your input data set. In addition, you can include the static SQL statements CONNECT, SET CONNECTION, or RELEASE. The maximum length for a statement is 32765 bytes. DSNTIAUL uses the SELECT statements to determine which tables to unload and dynamically executes all other statements except CONNECT, SET CONNECTION, and RELEASE. DSNTIAUL executes CONNECT, SET CONNECTION, and RELEASE statically to connect to remote locations.

If you do not specify the SQL parameter, your input data set must contain one or more single-line statements (without a semi-colon) that use the following syntax:

```
table or view name [WHERE conditions] [ORDER BY columns]
```

Each input statement must be a valid SQL SELECT statement with the clause SELECT * FROM omitted and with no ending semi-colon. DSNTIAUL generates a SELECT statement for each input statement by appending your input line to SELECT * FROM, then uses the result to determine which tables to unload. For this input format, the text for each table specification can be a maximum of 72 bytes and must not span multiple lines.

For both input formats, you can specify SELECT statements that join two or more tables or select specific columns from a table. If you specify columns, you will need to modify the LOAD statement that DSNTIAUL generates.

DSNTIAUL data sets:

Data set **Description**

SYSIN Input data set. See *DSNTIAUL parameters* for information on the contents of the input data.

You cannot enter comments in DSNTIAUL input.

The record length for the input data set must be at least 72 bytes. DSNTIAUL reads only the first 72 bytes of each record.

SYSPRINT Output data set. DSNTIAUL writes informational and error messages in this data set.

The record length for the SYSPRINT data set is 121 bytes.

SYSPUNCH Output data set. DSNTIAUL writes the LOAD utility control statements in this data set.

SYSREC*nn* Output data sets. The value *nn* ranges from 00 to 99. You can have a maximum of 100 output data sets for a single execution of DSNTIAUL. Each data set contains the data unloaded when DSNTIAUL processes a SELECT statement from the input data set. Therefore, the number of output data sets must match the number of SELECT statements (if you specify parameter SQL) or table specifications in your input data set.

Define all data sets as sequential data sets. You can specify the record length and block size of the SYSPUNCH and SYSREC*nn* data sets. The maximum record length for the SYSPUNCH and SYSREC*nn* data sets is 32760 bytes.

DSNTIAUL return codes:

Return code	Meaning
0	Successful completion.
4	An SQL statement received a warning code. If the SQL statement was a SELECT statement, DB2 did not perform the associated unload operation.
8	An SQL statement received an error code. If the SQL statement was a SELECT statement, DB2 did not perform the associated unload operation.
12	DSNTIAUL could not open a data set, an SQL statement returned a severe error code (-8*nn* or -9*nn*), or an error occurred in the SQL message formatting routine.

Examples of DSNTIAUL invocation: Suppose you want to unload the rows for department D01 from the project table. You can fit the table specification on one line, and you do not want to execute any non-SELECT statements, so you do not need the SQL parameter. Your invocation looks like this:

```
//UNLOAD  EXEC PGM=IKJEFT01,DYNAMNBR=20
//SYSTSPRT DD  SYSOUT=*
//SYSTSIN  DD  *
 DSN SYSTEM(DSN)
 RUN  PROGRAM(DSNTIAUL) PLAN(DSNTIB71) -
       LIB('DSN710.RUNLIB.LOAD')
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSREC00 DD DSN=DSN8UNLD.SYSREC00,
//            UNIT=SYSDA,SPACE=(32760,(1000,500)),DISP=(,CATLG),
//            VOL=SER=SCR03
//SYSPUNCH DD DSN=DSN8UNLD.SYSPUNCH,
//            UNIT=SYSDA,SPACE=(800,(15,15)),DISP=(,CATLG),
//            VOL=SER=SCR03,RECFM=FB,LRECL=120,BLKSIZE=1200
//SYSIN    DD *
DSN8710.PROJ WHERE DEPTNO='D01'
```

Figure 52. DSNTIAUL Invocation without the SQL parameter

If you want to obtain the LOAD utility control statements for loading rows into a table, but you do not want to unload the rows, you can set the data set names for the SYSREC*nn* data sets to DUMMY. For example, to obtain the utility control statements for loading rows into the department table, you invoke DSNTIAUL like this:

```
//UNLOAD  EXEC PGM=IKJEFT01,DYNAMNBR=20
//SYSTSPRT DD  SYSOUT=*
//SYSTSIN  DD  *
 DSN SYSTEM(DSN)
 RUN  PROGRAM(DSNTIAUL) PLAN(DSNTIB71) -
       LIB('DSN710.RUNLIB.LOAD')
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSREC00 DD DUMMY
//SYSPUNCH DD DSN=DSN8UNLD.SYSPUNCH,
//            UNIT=SYSDA,SPACE=(800,(15,15)),DISP=(,CATLG),
//            VOL=SER=SCR03,RECFM=FB,LRECL=120,BLKSIZE=1200
//SYSIN    DD *
DSN8710.DEPT
```

Figure 53. DSNTIAUL Invocation to obtain LOAD control statements

Now suppose that you also want to use DSNTIAUL to do these things:

- Unload all rows from the project table
- Unload only rows from the employee table for employees in departments with department numbers that begin with D, and order the unloaded rows by employee number
- Lock both tables in share mode before you unload them

For these activities, you must specify the SQL parameter when you run DSNTIAUL. Your DSNTIAUL invocation looks like this:

```
//UNLOAD  EXEC PGM=IKJEFT01,DYNAMNBR=20
//SYSTSPRT DD  SYSOUT=*
//SYSTSIN  DD  *
 DSN SYSTEM(DSN)
 RUN  PROGRAM(DSNTIAUL) PLAN(DSNTIB71) PARMS('SQL') -
       LIB('DSN710.RUNLIB.LOAD')
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSREC00 DD DSN=DSN8UNLD.SYSREC00,
//            UNIT=SYSDA,SPACE=(32760,(1000,500)),DISP=(,CATLG),
//            VOL=SER=SCR03
//SYSREC01 DD DSN=DSN8UNLD.SYSREC01,
//            UNIT=SYSDA,SPACE=(32760,(1000,500)),DISP=(,CATLG),
//            VOL=SER=SCR03
//SYSPUNCH DD DSN=DSN8UNLD.SYSPUNCH,
//            UNIT=SYSDA,SPACE=(800,(15,15)),DISP=(,CATLG),
//            VOL=SER=SCR03,RECFM=FB,LRECL=120,BLKSIZE=1200
//SYSIN    DD *
LOCK TABLE DSN8710.EMP IN SHARE MODE;
LOCK TABLE DSN8710.PROJ IN SHARE MODE;
SELECT * FROM DSN8710.PROJ;
SELECT * FROM DSN8710.EMP
  WHERE WORKDEPT LIKE 'D%'
  ORDER BY EMPNO;
```

Figure 54. DSNTIAUL Invocation with the SQL parameter

Running DSNTIAD

This section contains information that you need when you run DSNTIAD, including parameters, data sets, return codes, and invocation examples.

DSNTIAD parameters:

RC0

If you specify this parameter, DSNTIAD ends with return code 0, even if the program encounters SQL errors. If you do not specify RC0, DSNTIAD ends with a return code that reflects the severity of the errors that occur. Without RC0, DSNTIAD terminates if more than 10 SQL errors occur during a single execution.

SQLTERM(*termchar***)**

Specify this parameter to indicate the character that you use to end each SQL statement. You can use any special character *except* one of those listed in Table 139. SQLTERM(;) is the default.

Table 139. Invalid special characters for the SQL terminator

Name	Character	Hexadecimal representation
blank		X'40'
comma	,	X'5E'
double quote	"	X'7F'
left parenthesis	(	X'4D'
right parenthesis	)	X'5D'
single quote	'	X'7D'
underscore	_	X'6D'

Use a character other than a semicolon if you plan to execute a statement that contains embedded semicolons. For example, suppose you specify the parameter SQLTERM(#) to indicate that the character # is the statement terminator. Then a CREATE TRIGGER statement with embedded semicolons looks like this:

```
CREATE TRIGGER NEW_HIRE
  AFTER INSERT ON EMP
  FOR EACH ROW MODE DB2SQL
  BEGIN ATOMIC
    UPDATE COMPANY_STATS SET NBEMP = NBEMP + 1;
  END#
```

Be careful to choose a character for the statement terminator that is not used within the statement.

DSNTIAD data sets:

Data Set	Description
SYSIN	Input data set. In this data set, you can enter any number of non-SELECT SQL statements, each terminated with a semi-colon. A statement can span multiple lines, but DSNTIAD reads only the first 72 bytes of each line. You cannot enter comments in DSNTIAD input.
SYSPRINT	Output data set. DSNTIAD writes informational and error messages in this data set. DSNTIAD sets the record length of this data set to 121 and the block size to 1210.

Define all data sets as sequential data sets.

DSNTIAD return codes:

Return code	Meaning
0	Successful completion, or the user specified parameter RC0.
4	An SQL statement received a warning code.
8	An SQL statement received an error code.
12	DSNTIAD could not open a data set, the length of an SQL statement was more than 32 760 bytes, an SQL statement returned a severe error code (-8*nn* or -9*nn*), or an error occurred in the SQL message formatting routine.

Example of DSNTIAD invocation: Suppose you want to execute 20 UPDATE statements, and you do not want DSNTIAD to terminate if more than 10 errors occur. Your invocation looks like this:

```
//RUNTIAD EXEC PGM=IKJEFT01,DYNAMNBR=20
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *
 DSN SYSTEM(DSN)
 RUN  PROGRAM(DSNTIAD) PLAN(DSNTIA71) PARMS('RC0') -
       LIB('DSN710.RUNLIB.LOAD')
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN    DD *
UPDATE DSN8710.PROJ SET DEPTNO='J01' WHERE DEPTNO='A01';
UPDATE DSN8710.PROJ SET DEPTNO='J02' WHERE DEPTNO='A02';
:

UPDATE DSN8710.PROJ SET DEPTNO='J20' WHERE DEPTNO='A20';
```

Figure 55. DSNTIAD Invocation with the RC0 Parameter

Running DSNTEP2

This section contains information that you need when you run DSNTEP2, including parameters, data sets, return codes, and invocation examples.

DSNTEP2 parameters:

Parameter
Description

ALIGN(MID) or ALIGN(LHS)
If you want your DSNTEP2 output centered, specify ALIGN(MID). If you want the output left-aligned, choose ALIGN(LHS). The default is ALIGN(MID).

NOMIXED or MIXED
If your input to DSNTEP2 contains any DBCS characters, specify MIXED. If your input contains no DBCS characters, specify NOMIXED. The default is NOMIXED.

SQLTERM(*termchar***)**
Specify this parameter to indicate the character that you use to end each SQL statement. You can use any character *except* one of those listed in Table 139 on page 689. SQLTERM(;) is the default.

Use a character other than a semicolon if you plan to execute a statement that contains embedded semicolons. For example, suppose you specify the parameter SQLTERM(#) to indicate that the character # is the statement terminator. Then a CREATE TRIGGER statement with embedded semicolons looks like this:

```
CREATE TRIGGER NEW_HIRE
  AFTER INSERT ON EMP
  FOR EACH ROW MODE DB2SQL
  BEGIN ATOMIC
    UPDATE COMPANY_STATS SET NBEMP = NBEMP + 1;
  END#
```

Be careful to choose a character for the statement terminator that is not used within the statement.

If you want to change the SQL terminator within a series of SQL statements, you can use the --#SET TERMINATOR control statement. For example, suppose that you have an existing set of SQL statements to which you want to

add a CREATE TRIGGER statement that has embedded semicolons. You can use the default SQLTERM value, which is a semicolon, for all of the existing SQL statements. Before you execute the CREATE TRIGGER statement, include the --#SET TERMINATOR # control statement to change the SQL terminator to the character #:

```
SELECT * FROM DEPT;
SELECT * FROM ACT;
SELECT * FROM EMPPROJACT;
SELECT * FROM PROJ;
SELECT * FROM PROJACT;
--#SET TERMINATOR #
CREATE TRIGGER NEW_HIRE
  AFTER INSERT ON EMP
  FOR EACH ROW MODE DB2SQL
  BEGIN ATOMIC
    UPDATE COMPANY_STATS SET NBEMP = NBEMP + 1;
  END#
```

See the discussion of the SYSIN data set for more information on the --#SET control statement.

DSNTEP2 data sets:

Data Set	Description
SYSIN	Input data set. In this data set, you can enter any number of SQL statements, each terminated with a semi-colon. A statement can span multiple lines, but DSNTEP2 reads only the first 72 bytes of each line. You can enter comments in DSNTEP2 input with an asterisk (*) in column 1 or two hyphens (--) anywhere on a line. Text that follows the asterisk is considered to be comment text. Text that follows two hyphens can be comment text or a control statement. Comments and control statements cannot span lines. You can enter a number of control statements in the DSNTEP2 input data set. Those control statements are of the form `--#SET control-option value` The control options are: **TERMINATOR** The SQL statement terminator. *value* is any single-byte character other than one of those listed in Table 139 on page 689. The default is the value of the SQLTERM parameter. **ROWS_FETCH** The number of rows to be fetched from the result table. *value* is a numeric literal between -1 and the number of rows in the result table. -1 means that all rows are to be fetched. The default is -1. **ROWS_OUT** The number of fetched rows to be sent to the output data set. *value* is a numeric literal between -1 and the number of fetched rows. -1 means that all fetched rows are to be sent to the output data set. The default is -1.

SYSPRINT Output data set. DSNTEP2 writes informational and error messages in this data set. DSNTEP2 writes output records of no more than 133 bytes.

Define all data sets as sequential data sets.

DSNTEP2 return codes:

Return code	Meaning
0	Successful completion.
4	An SQL statement received a warning code.
8	An SQL statement received an error code.
12	The length of an SQL statement was more than 32 760 bytes, an SQL statement returned a severe error code (-8*nn* or -9*nn*), or an error occurred in the SQL message formatting routine.

Example of DSNTEP2 invocation: Suppose you want to use DSNTEP2 to execute SQL SELECT statements that might contain DBCS characters. You also want your output left-aligned. Your invocation looks like this:

```
//RUNTEP2 EXEC PGM=IKJEFT01,DYNAMNBR=20
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *
 DSN SYSTEM(DSN)
 RUN  PROGRAM(DSNTEP2) PLAN(DSNTEP71) PARMS('/ALIGN(LHS) MIXED') -
       LIB('DSN710.RUNLIB.LOAD')
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN    DD *
SELECT * FROM DSN8710.PROJ;
```

Figure 56. DSNTEP2 Invocation with the ALIGN(LHS) and MIXED parameters

Appendix E. DB2 utilities packaging

DB2 Version 7 provides utilities in various packages. Core utilities are included with DB2. Other utilities are available in three separate products. All DB2 utilities operate on catalog, directory and sample objects: without requiring any additional products.

The core utilities include:

- CATMAINT
- DIAGNOSE
- LISTDEF
- OPTIONS
- QUIESCE
- REPAIR
- REPORT
- TEMPLATE
- all standalone utilities

The DB2 utility products are:

- 5655-E63: DB2 Operational Utilities (FMID JDB771K)
  - COPY
  - EXEC SQL
  - LOAD
  - REBUILD INDEX
  - RECOVER
  - REORG INDEX
  - REORG TABLESPACE
  - RUNSTATS
  - STOSPACE
  - UNLOAD
- 5655-E62: DB2 Diagnostic and Recovery Utilities (FMID JDB771M)
  - CHECK DATA
  - CHECK INDEX
  - CHECK LOB
  - COPY
  - COPYTOCOPY
  - MERGECOPY
  - MODIFY RECOVERY
  - MODIFY STATISTICS
  - REBUILD INDEX
  - RECOVER
- 5697-E98: DB2 Utilities Suite (FMIDs JDB771K and JDB771M)
  - CHECK DATA
  - CHECK INDEX
  - CHECK LOB
  - COPY
  - COPYTOCOPY
  - EXEC SQL
  - LOAD
  - MERGECOPY
  - MODIFY RECOVERY
  - MODIFY STATISTICS
  - REBUILD INDEX
  - RECOVER
  - REORG INDEX

© Copyright IBM Corp. 1983, 2001

- REORG TABLESPACE
- RUNSTATS
- STOSPACE
- UNLOAD

SMP/E jobs for DB2 utility products

To load the DB2 utility products, use System Modification Program Extended (SMP/E). SMP/E processes the installation tapes or cartridges and creates DB2 distribution target libraries.

Several jobs are provided that invoke SMP/E. These jobs are on the tape or cartridge you received with the utility product. The job prologues in these jobs contain directions on how to tailor the job for your site. Follow these directions carefully to ensure that your DB2 for OS/390 and z/OS SMP/E process functions correctly. To copy the jobs from the tapes submit the copy job listed in the DB2 Program Directory.

Installing DB2 Operational Utilities

The SMP/E RECEIVE job, DSNRECVK, loads the DB2 Operational Utilities program modules, macros, and procedures into temporary data sets (SMPTLIBs). If this job fails or abends, correct the problem and rerun the job. Use job DSNRECV1 described in *DB2 Installation Guide* as a guide to help you with this job.

The SMP/E APPLY job, DSNAPPLK, copies and link-edits the DB2 program modules, macros, and procedures into the DB2 target libraries. Use job DSNAPPL1 described in *DB2 Installation Guide* as a guide to help you with this job.

The SMP/E ACCEPT job, DSNACCPK, copies the DB2 Operational Utility program modules, macros, and procedures into the DB2 distribution libraries. Use job DSNACEP1 described in *DB2 Installation Guide* as a guide to help you with this job.

Installing DB2 Diagnostic & Recovery Utilities

The SMP/E RECEIVE job, DSNRECVM, loads the DB2 Diagnostic and Recovery Utilities program modules, macros, and procedures into temporary data sets (SMPTLIBs). If this job fails or abends, correct the problem and rerun the job. Use job DSNRECV1 described in *DB2 Installation Guide* as a guide to help you with this job.

The SMP/E APPLY job, DSNAPPLM, copies and link-edits the DB2 Diagnostic and Recovery Utilities program modules, macros, and procedures into the DB2 target libraries. Use job DSNAPPL1 described in *DB2 Installation Guide* as a guide to help you with this job.

The SMP/E ACCEPT job, DSNACCPM, copies the DB2 Diagnostic and Recovery Utilities program modules, macros, and procedures into the DB2 distribution libraries. Use job DSNACEP1 described in *DB2 Installation Guide* as a guide to help you with this job.

Installing DB2 Utilities Suite

The SMP/E RECEIVE job, DSNRECVS, loads the DB2 Diagnostic and Recovery Utilities program modules, macros, and procedures into temporary data sets (SMPTLIBs). The SMP/E RECEIVE job, DSNRECVK, loads the DB2 Operational

Utilities program modules, macros, and procedures into temporary data sets (SMPTLIBs). If these jobs fail or abend, correct the problem and rerun the jobs. Use job DSNRECV1 described in *DB2 Installation Guide* as a guide to help you with the RECEIVE job.

The SMP/E APPLY job, DSNAPPLS, copies and link-edits both the DB2 Diagnostic and Recovery Utilities program modules, macros, and procedures; and the the DB2 Operational Utilities program modules, macros, and procedures into the DB2 target libraries. Use job DSNAPPL1 described in *DB2 Installation Guide* as a guide to help you with the APPLY job.

The SMP/E ACCEPT job, DSNACCPS, copies the DB2 Diagnostic and Recovery Utilities program modules, macros, and procedures; and the DB2 Operational Utility program modules, macros, and procedures into the DB2 distribution libraries. Use job DSNACEP1 described in *DB2 Installation Guide* as a guide to help you with the ACCEPT job.

How to operate in a mixed-release data sharing environment

The utilities batch module DSNUTILB, is split into multiple parts: a release-independent module called DSNUTILB and multiple release-dependent modules, DSNUT710 and the utility-dependent load modules listed in Table 140. To operate in a mixed-release data sharing environment, you must have the release-dependent modules from both releases and all applicable utility-dependent modules available to the utility jobs that operate across the data sharing group. The procedure for sharing utility modules is explained in Chapter 4 of *DB2 Data Sharing: Planning and Administration*. Use the information in Table 140 and the procedures outlined in Chapter 4 of *DB2 Data Sharing: Planning and Administration* to implement a mixed-release data sharing environment.

With Version 7, there are separate load modules and aliases for each utility. For a cross reference between the utility name, alias name and load module names, see Table 140. Each utility product contains unique load module names and some utilities are available in multiple utility products. Therefore some utilities have multiple load module names because the utility ships with different products. The core utility load modules use DSNU7CL*x* naming convention. The Operational utilities use DSNU7OL*x* names. The Diagnostic and Recovery utilities us DSNU7RL*x* names.

Table 140. Relationship between utility names, aliases and utility load modules

Utility name	Alias name	Load module name
CATMAINT	DSNU71AA	DSBU7CLA
CHECK	DSNU71AB	DSNU7RLB
COPY	DSNU71AC	DSNU7OLC, DSNU7RLC
COPYTOCOPY	DSNU71AT	DSNU7RLT
DIAGNOSE	DSNU71AD	DSNU7CLD
EXEC SQL	DSNU71AU	DSNU7OLU
LISTDEF	DSNU71AE	DSNU7CLE
LOAD	DSNU71AF	DSNU7OLF
MERGECOPY	DSNU71AG	DSNU7RLG
MODIFY RECOVERY and MODIFY STATISTICS	DSNU71AH	DSNU7RLH

Table 140. Relationship between utility names, aliases and utility load modules (continued)

Utility name	Alias name	Load module name
OPTIONS	DSNU71AI	DSNU7CLI
QUIESCE	DSNU71AJ	DSNU7CLJ
REBUILD INDEX	DSNU71AK	DSNU7OLK, DSNU7RLK
RECOVER	DSNU71AL	DSNU7OLL, DSNU7RLL
REORG INDEX and REORG TABLESPACE	DSNU71AM	DSNU7OLM
REPAIR	DSNU71AN	DSNU7CLN
REPORT	DSNU71AO	DSNU7CLO
RUNSTATS	DSNU71AP	DSNU7OLP
STOSPACE	DSNU71AQ	DSNU7OLQ
TEMPLATE	DSNU71AR	DSNU7CLR
UNLOAD	DSNU71AS	DSNU7OLS

Appendix F. Real-time statistics tables

The information under this heading is Product-sensitive Programming Interface and Associated Guidance Information, as defined in “Notices” on page 725.

DB2 collects statistics that you can use to determine when you need to perform certain maintenance functions on your table spaces and index spaces.

DB2 collects the statistics in real time. You create tables into which DB2 periodically writes the statistics. You can then write applications that query the statistics and help you decide when to run REORG, RUNSTATS, or COPY, or enlarge your data sets. For information on a DB2-supplied stored procedure that queries the real-time statistics tables, see “The DB2 real-time statistics stored procedure” on page 658.

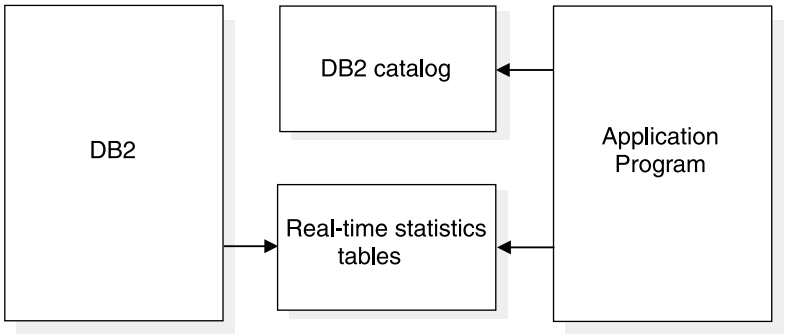

Figure 57. Real-Time Statistics overview

The following sections provide detailed information about the real-time statistics tables:

- “Setting up your system for real-time statistics”
- “Contents of the real-time statistics tables” on page 701
- “Operating with real-time statistics” on page 713

Setting up your system for real-time statistics

DB2 always generates in-memory statistics for each table space and index space in your system. For partitioned spaces, DB2 generates information for each partition. However, you need to perform the following steps before DB2 externalizes the statistics to DB2 tables:

1. Create the real-time statistics objects. See “Creating and altering the real-time statistics objects”.
2. Set the interval for writing statistics. See “Setting the interval for writing real-time statistics” on page 700.
3. Start the real-time statistics database. See “Starting the real-time statistics database” on page 701.

Creating and altering the real-time statistics objects

You need to create a database, table space, tables and indexes for the real-time statistics. Those objects are listed in Table 141 on page 700. Use the SQL statements in member DSNTESS of data set DSN710.SDSNSAMP as a model for creating the real-time statistics objects. You can create these objects in user-managed or DB2-managed data sets.

© Copyright IBM Corp. 1983, 2001

Restrictions on changing the provided definitions for the real-time statistics objects: You can change most of the attributes in the provided definitions of the real-time statistics objects. However, you cannot change the following items:

- Object names

  You must use the names that are specified in DSNTESS for the database, table space, tables, indexes, and table columns.
- The CCSID parameter on the CREATE DATABASE, CREATE TABLESPACE, and CREATE TABLE statements

  The CCSID must be EBCDIC.
- Number of columns or column definitions

  You cannot add table columns or modify column definitions.

Before you can alter an object in the real-time statistics database, you must stop the database. Otherwise, you receive an SQL error.

Table 141. DB2 objects for storing real-time statistics

Object name	Description
DSNRTSDB	Database for real-time statistics objects
DSNRTSTS	Table space for real-time statistics objects
SYSIBM.TABLESPACESTATS	Table for statistics on table spaces and table space partitions
SYSIBM.INDEXSPACESTATS	Table for statistics on index spaces and index space partitions
SYSIBM.TABLESPACESTATS_IX	Unique Index on SYSIBM.TABLESPACESTATS (columns DBID, PSID, and PARTITION)
SYSIBM.INDEXSPACESTATS_IX	Unique Index on SYSIBM.INDEXSPACESTATS (columns DBID, PSID, and PARTITION)

To create the real-time statistics objects, you need the authority to create tables and indexes on behalf of the SYSIBM authorization ID.

DB2 inserts one row in the table for each partition or non-partitioned table space or index space. You therefore need to calculate the amount of disk space that you need for the real-time statistics tables based on the current number of table spaces and indexes in your subsystem.

To determine the amount of storage that you need for the real-time statistics when they are in memory, estimate the peak number of objects that might be updated concurrently, and multiply that total by the amount of in-memory space that DB2 uses for each object (152 bytes):

```
Amount of Storage in bytes = Maximum concurrent objects updated * 152 bytes
```

Recommendation: Place the statistics indexes and tables in their own buffer pool. When the statistics pages are in memory, the speed at which in-memory statistics are written to the tables improves.

Setting the interval for writing real-time statistics

You can set the interval for writing real-time statistics when you install DB2 and update that interval online. The installation field is REAL TIME STATS on panel DSNTIPO. The default interval is 30 minutes. To update the interval, modify system parameter STATSINT.

In a data sharing environment, each member has its own interval for writing real-time statistics.

Starting the real-time statistics database

After you create the real-time statistics database, DB2 puts it into a stopped state. After you create all the objects in the database, you need to issue START DATABASE(DSNRTSDB) to explicitly start the database.

You must start the database in read-write mode to make it possible for DB2 to externalize real-time statistics. See “When DB2 externalizes real-time statistics” on page 713 for information on when DB2 externalizes the statistics.

Contents of the real-time statistics tables

The SYSIBM.TABLESPACESTATS table contains statistics information on table spaces and table space partitions. The SYSIBM.INDEXSPACESTATS table contains statistics information on index spaces and index space partitions.

Table 142 describes the columns of the TABLESPACESTATS table and explains how you can use them in deciding when to run REORG, RUNSTATS, or COPY.

Table 142. Descriptions of columns in the TABLESPACESTATS table

Column name	Data type	Description
DBNAME	CHAR(8) NOT NULL	The name of the database. This column is used to map a database to its statistics.
NAME	CHAR(8) NOT NULL	The name of the table space. This column is used to map a table space to its statistics.
PARTITION	SMALLINT NOT NULL	The data set number within the table space. This column is used to map a data set number in a table space to its statistics. For partitioned table spaces, this value corresponds to the partition number for a single partition. For nonpartitioned table spaces, this value is 0.
DBID	SMALLINT NOT NULL	The internal identifier of the database. This column is used to map a DBID to its statistics.
PSID	SMALLINT NOT NULL	The internal identifier of the table space page set descriptor. This column is used to map a PSID to its statistics.

Table 142. Descriptions of columns in the TABLESPACESTATS table (continued)

Column name	Data type	Description
UPDATESTATSTIME	TIMESTAMP NOT NULL WITH DEFAULT	The timestamp when the row was inserted or last updated. This column is updated with the current timestamp when a row in the TABLESPACESTATS table is inserted or updated. You can use this column in several ways: • To determine the actions that caused the latest change to the table. Do this by selecting any of the timestamp columns and comparing them to the UPDATESTATSTIME column. • To determine whether an analysis of data is needed. This determination might be based on a given time interval, or on a combination of the time interval and the amount of activity. For example, suppose you want to analyze statistics for the last seven days. To determine whether there has been any activity in the past seven days, check whether the difference between the current date and the UPDATESTATSTIME value is less than or equal to seven: `(JULIAN_DAY(CURRENT DATE)-JULIAN_DAY(UPDATESTATSTIME))<= 7`
TOTALROWS	FLOAT	The number of rows or LOBs in the table space or partition. If the table space contains more than one table, this value is the sum of all rows in all tables. A null value means that the number of rows is unknown, or REORG or LOAD has never been run. Use this value with the value of any column that contains a number of affected rows to determine the percentage of rows that are affected by a particular action.
NACTIVE	INTEGER	The number of active pages in the table space or partition. A null value means the number of active pages is unknown. This value is equivalent to the number of preformatted pages. For multi-piece table spaces, this value is the total number of preformatted pages in all data sets. Use this value with the value of any column that contains a number of affected pages to determine the percentage of pages that are affected by a particular action. For example, suppose that your site's maintenance policies require that COPY is run after 20 per cent of the pages in a table space have changed. To determine if a COPY might be required, calculate the ratio of updated pages since the last COPY to the total number of active pages. If the percentage is greater than 20, you need to run COPY: `((COPYUPDATEDPAGES*100)/NACTIVE)>20`
SPACE	INTEGER	The amount of space that is allocated to the table space or partition, in kilobytes. For multi-piece linear page sets, this value is the amount of space in all data sets. A null value means the amount of space is unknown. Use this value to monitor growth and validate design assumptions.

Table 142. Descriptions of columns in the TABLESPACESTATS table (continued)

#	Column name	Data type	Description
#	EXTENTS	SMALLINT	The number of physical extents in the table space or partition. # For multi-piece linear page sets, this value is the number of extents # for the last data set. A null value means the number of extents is # unknown. # Use this value to determine: # • When the primary or secondary allocation value for a table # space or partition needs to be altered. # • When you are approaching the maximum number of extents and # risking extend failures.
#	LOADRLASTTIME	TIMESTAMP	The timestamp of the last LOAD REPLACE on the table space or # partition. # A null value means LOAD REPLACE has never been run on the # table space or partition, or the timestamp of the last LOAD # REPLACE is unknown. # You can compare this timestamp to the timestamp of the last # COPY on the same object to determine when a COPY is needed. If # the date of the last LOAD REPLACE is more recent than the last # COPY, you might need to run COPY: # `(JULIAN_DAY(LOADRLASTTIME)>JULIAN_DAY(COPYLASTTIME))`
#	REORGLASTTIME	TIMESTAMP	The timestamp of the last REORG on the table space or partition. # A null value means REORG has never been run on the table space # or partition, or the timestamp of the last REORG is unknown. # You can compare this timestamp to the timestamp of the last # COPY on the same object to determine when a COPY is needed. If # the date of the last REORG is more recent than the last COPY, you # might need to run COPY: # `(JULIAN_DAY(REORGLASTTIME)>JULIAN_DAY(COPYLASTTIME))`
#	REORGINSERTS	INTEGER	The number of records or LOBs that have been inserted since the # last REORG or LOAD REPLACE on the table space or partition. # A null value means that the number of inserted records or LOBs is # unknown.
#	REORGDELETES	INTEGER	The number of records or LOBs that have been deleted since the # last REORG or LOAD REPLACE on the table space or partition. # A null value means that the number of deleted records or LOBs is # unknown.

Table 142. Descriptions of columns in the TABLESPACESTATS table (continued)

Column name	Data type	Description
REORGUPDATES	INTEGER	The number of rows that have been updated since the last REORG or LOAD REPLACE on the table space or partition. This value does not include LOB updates because LOB updates are really deletions followed by insertions. A null value means that the number of updated rows is unknown. This value can be used with REORGDELETES and REORGINSERTS to determine if a REORG is necessary. For example, suppose that your site's maintenance policies require that REORG is run after 20 per cent of the rows in a table space have changed. To determine if a REORG is required, calculate the sum of updated, inserted, and deleted rows since the last REORG. Then calculate the ratio of that sum to the total number of rows. If the percentage is greater than 20, you might need to run REORG: `(((REORGINSERTS+REORGDELETES+REORGUPDATES)*100)/TOTALROWS)>20`
REORGDISORGLOB	INTEGER	The number of LOBs that were inserted since the last REORG or LOAD REPLACE that are not perfectly chunked. A LOB is perfectly chunked if the allocated pages are in the minimum number of chunks. A null value means that the number of imperfectly chunked LOBs is unknown. Use this value to determine whether you need to run REORG. For example, you might want to run REORG if the ratio of REORGDISORGLOB to the total number of LOBs is greater than 10 per cent: `((REORGDISORGLOB*100)/TOTALROWS)>10`
REORGUNCLUSTINS	INTEGER	The number of records that were inserted since the last REORG or LOAD REPLACE. that are not well-clustered with respect to the clustering index. A record is well-clustered if the record is inserted into a page that is within 16 pages of the ideal candidate page. The clustering index determines the ideal candidate page. A null value means that the number of badly-clustered pages is unknown. You can use this value to determine whether you need to run REORG. For example, you might want to run REORG if the following comparison is true: `((REORGUNCLUSTINS*100)/TOTALROWS)>10`
REORGMASSDELETE	INTEGER	The number of mass deletes from a segmented or LOB table space, or the number of dropped tables from a segmented table space, since the last REORG or LOAD REPLACE. A null value means that the number of mass deletes is unknown. If this value is non-zero, a REORG might be necessary.
REORGNEARINDREF	INTEGER	The number of overflow records that were created since the last REORG or LOAD REPLACE and were relocated near the pointer record. For nonsegmented table spaces, a page is near the present page if the two page numbers differ by 16 or less. For segmented table spaces, a page is near the present page if the two page numbers differ by SEGSIZE*2 or less. A null value means that the number of overflow records near the pointer record is unknown.

Table 142. Descriptions of columns in the TABLESPACESTATS table (continued)

Column name	Data type	Description

REORGFARINDEF	INTEGER	The number of overflow records that were created since the last REORG or LOAD REPLACE and were relocated far from the pointer record.
		
		For nonsegmented table spaces, a page is far from the present page if the two page numbers differ by more than 16. For segmented table spaces, a page is far from the present page if the two page numbers differ by at least (SEGSIZE*2)+1. A null value means that the number of overflow records far from the pointer record is unknown.
		
		For example, in a non-data sharing environment, you might run REORG if the following comparison is true:
		`(((REORGNEARINDREF+REORGFARINDREF)*100)/TOTALROWS)>10`
		
		In a data sharing environment, you might run REORG if the following comparison is true:
		`(((REORGNEARINDREF+REORGFARINDREF)*100)/TOTALROWS)>5`

STATSLASTTIME	TIMESTAMP	The timestamp of the last RUNSTATS on the table space or partition. A null value means RUNSTATS has never been run on the table space or partition, or the timestamp of the last RUNSTATS is unknown. You can compare this timestamp to the timestamp of the last REORG on the same object to determine when RUNSTATS is needed. If the date of the last REORG is more recent than the last RUNSTATS, you might need to run RUNSTATS: `(JULIAN_DAY(REORGLASTTIME)>JULIAN_DAY(STATSLASTTIME))`
STATSINSERTS	INTEGER	The number of records or LOBs that have been inserted since the last RUNSTATS on the table space or partition. A null value means that the number of inserted records or LOBs is unknown.
STATSDELETES	INTEGER	The number of records or LOBs that have been deleted since the last RUNSTATS on the table space or partition. A null value means that the number of deleted records or LOBs is unknown.
STATSUPDATES	INTEGER	The number of rows that have been updated since the last RUNSTATS on the table space or partition. This value does not include LOB updates because LOB updates are really deletions followed by insertions. A null value means that the number of updated rows is unknown. This value can be used with STATSDELETES and STATSINSERTS to determine if RUNSTATS is necessary. For example, suppose that your site's maintenance policies require that RUNSTATS is run after 20 per cent of the rows in a table space have changed. To determine if RUNSTATS is required, calculate the sum of updated, inserted, and deleted rows since the last RUNSTATS. Then calculate the ratio of that sum to the total number of rows. If the percentage is greater than 20, you need to run RUNSTATS: `(((STATSINSERTS+STATSDELETES+STATSUPDATES)*100)/TOTALROWS)>20`

Table 142. Descriptions of columns in the TABLESPACESTATS table (continued)

Column name	Data type	Description
STATSMASSDELETE	INTEGER	The number of mass deletes from a segmented or LOB table space, or the number of dropped tables from a segmented table space, since the last RUNSTATS. A null value means that the number of mass deletes is unknown. If this value is non-zero, RUNSTATS might be necessary.
COPYLASTTIME	TIMESTAMP	The timestamp of the last full or incremental image copy on the table space or partition. A null value means COPY has never been run on the table space or partition, or the timestamp of the last full image copy is unknown. You can compare this timestamp to the timestamp of the last REORG on the same object to determine when a COPY is needed. If the date of the last REORG is more recent than the last COPY, you might need to run COPY: `(JULIAN_DAY(REORGRLASTTIME)>JULIAN_DAY(COPYLASTTIME))`

COPYUPDATEDPAGES	INTEGER	The number of distinct pages that have been updated since the last COPY.
		
		A null value means that the number of updated pages is unknown.
		
		You can compare this value to the total number of pages to determine when a COPY is needed.
		
		For example, you might want to take an incremental image copy when one percent of the pages have changed:
		`((COPYUPDATEDPAGES*100)/NACTIVE)>1`
		
		You might want to take a full image copy when 20 percent of the pages have changed:
		`((COPYUPDATEDPAGES*100)/NACTIVE)>20`

COPYCHANGES	INTEGER	The number of insert, delete, and update operations since the last COPY.
		
		A null value means that the number of insert, delete, or update operations is unknown.
		
		This number indicates the approximate number of log records that DB2 processes to recover to the current state.
		
		For example, you might want to take an incremental image copy when DB2 processes more than one percent of the rows from the logs:
		`((COPYCHANGES*100)/TOTALROWS)>1`
		
		You might want to take a full image copy when DB2 processes more than 10 percent of the rows from the logs:
		`((COPYCHANGES*100)/TOTALROWS)>10`

Table 142. Descriptions of columns in the TABLESPACESTATS table (continued)

Column name	Data type	Description
COPYUPDATELRSN	CHAR(6) FOR BIT DATA	The LRSN or RBA of the first update after the last COPY. A null value means that the LRSN or RBA is unknown. Consider running COPY if this value is not in the active logs. To determine the oldest LRSN or RBA in the active logs, use the Print Log Map utility (DSNJU004).
COPYUPDATETIME	TIMESTAMP	The timestamp of the first update after the last COPY. A null value means that the timestamp is unknown. This value has a similar purpose to COPYUPDATELRSN.

Table 143 describes the columns of the INDEXSPACESTATS table and explains how you can use them in deciding when to run REORG, RUNSTATS, or COPY.

Table 143. Descriptions of columns in the INDEXSPACESTATS table

Column name	Data type	Description
DBNAME	CHAR(8) NOT NULL	The name of the database. This column is used to map a database to its statistics.
NAME	CHAR(8) NOT NULL	The name of the index space. This column is used to map an index space to its statistics.
PARTITION	SMALLINT NOT NULL	The data set number within the index space. This column is used to map a data set number in an index space to its statistics. For partitioned index spaces, this value corresponds to the partition number for a single partition. For nonpartitioned index spaces, this value is 0.
DBID	SMALLINT NOT NULL	The internal identifier of the database. This column is used to map a DBID to its statistics.
ISOBID	SMALLINT NOT NULL	The internal identifier of the index space page set descriptor. This column is used to map an ISOBID to its statistics.
PSID	SMALLINT NOT NULL	The internal identifier of the table space page set descriptor for the table space on which the index that is represented by this row is created. This column is used to map a PSID to the statistics for the associated index.

Table 143. Descriptions of columns in the INDEXSPACESTATS table (continued)

Column name	Data type	Description
UPDATESTATSTIME	TIMESTAMP NOT NULL WITH DEFAULT	The timestamp when the row was inserted or last updated. This column is updated with the current timestamp when a row in the INDEXSPACESTATS table is inserted or updated. You can use this column in several ways: • To determine the actions that caused the latest change to the INDEXSPACESTATS table. Do this by selecting any of the timestamp columns and comparing them to the UPDATESTATSTIME column. • To determine whether an analysis of data is needed. This determination might be based on a given time interval, or on a combination of the time interval and the amount of activity. For example, suppose you want to analyze statistics for the last seven days. To determine whether there has been any activity in the past seven days, check whether the difference between the current date and the UPDATESTATSTIME value is less than or equal to seven: `(JULIAN_DAY(CURRENT DATE)-JULIAN_DAY(UPDATESTATSTIME))<= 7`
TOTALENTRIES	FLOAT	The number of entries, including duplicate entries, in the index space or partition. A null value means that the number of entries is unknown, or REORG, LOAD, or REBUILD has never been run. Use this value with the value of any column that contains a number of affected index entries to determine the percentage of index entries that are affected by a particular action.
NLEVELS	SMALLINT	The number of levels in the index tree. A null value means that the number of levels is unknown.
NACTIVE	INTEGER	The number of active pages in the index space or partition. A null value means the number of active pages is unknown. This value is equivalent to the number of preformatted pages. Use this value with the value of any column that contains a number of affected pages to determine the percentage of pages that are affected by a particular action. For example, suppose that your site's maintenance policies require that COPY is run after 20 per cent of the pages in an index space have changed. To determine if a COPY is required, calculate the ratio of updated pages since the last COPY to the total number of active pages. If the percentage is greater than 20, you need to run COPY: `((COPYUPDATEDPAGES*100)/NACTIVE)>20`
SPACE	INTEGER	The amount of space that is allocated to the index space or partition, in kilobytes. For multi-piece linear page sets, this value is the amount of space in all data sets. A null value means the amount of space is unknown. Use this value to monitor growth and validate design assumptions.

Table 143. Descriptions of columns in the INDEXSPACESTATS table (continued)

# Column name	Data type	Description
# EXTENTS	SMALLINT	The number of physical extents in the index space or partition. # For multi-piece linear page sets, this value is the number of extents for the last data set. A null value means the number of extents is unknown. # Use this value to determine: # • When the primary allocation value for an index space or partition needs to be altered. # • When you are approaching the maximum number of extents and risking extend failures.
# LOADRLASTTIME	TIMESTAMP	The timestamp of the last LOAD REPLACE on the index space or partition. # A null value means that the timestamp of the last LOAD REPLACE is unknown. # If COPY YES was specified when the index was created (the value of COPY is Y in SYSIBM.SYSINDEXES), you can compare this timestamp to the timestamp of the last COPY on the same object to determine when a COPY is needed. If the date of the last LOAD REPLACE is more recent than the last COPY, you might need to run COPY: # `(JULIAN_DAY(LOADRLASTTIME)>JULIAN_DAY(COPYLASTTIME))`
# REBUILDLASTTIME	TIMESTAMP	The timestamp of the last REBUILD INDEX on the index space or partition. # A null value means the timestamp of the last REBUILD INDEX is unknown. # If COPY YES was specified when the index was created (the value of COPY is Y in SYSIBM.SYSINDEXES), you can compare this timestamp to the timestamp of the last COPY on the same object to determine when a COPY is needed. If the date of the last REBUILD INDEX is more recent than the last COPY, you might need to run COPY: # `(JULIAN_DAY(REBUILDLASTTIME)>JULIAN_DAY(COPYLASTTIME))`
# REORGLASTTIME	TIMESTAMP	The timestamp of the last REORG INDEX on the index space or partition. # A null value means the timestamp of the last REORG INDEX is unknown. # If COPY YES was specified when the index was created (the value of COPY is Y in SYSIBM.SYSINDEXES), you can compare this timestamp to the timestamp of the last COPY on the same object to determine when a COPY is needed. If the date of the last REORG INDEX is more recent than the last COPY, you might need to run COPY: # `(JULIAN_DAY(REORGLASTTIME)>JULIAN_DAY(COPYLASTTIME))`
# REORGINSERTS	INTEGER	The number of index entries that have been inserted since the last REORG, REBUILD INDEX or LOAD REPLACE on the index space or partition. # A null value means that the number of inserted index entries is unknown.

Table 143. Descriptions of columns in the INDEXSPACESTATS table (continued)

# Column name	Data type	Description
# REORGDELETES	INTEGER	The number of index entries that have been deleted since the last REORG, REBUILD INDEX, or LOAD REPLACE on the index space or partition. A null value means that the number of deleted index entries is unknown. This value can be used with REORGINSERTS to determine if a REORG is necessary. For example, suppose that your site's maintenance policies require that REORG is run after 20 per cent of the index entries have changed. To determine if a REORG is required, calculate the sum of inserted and deleted rows since the last REORG. Then calculate the ratio of that sum to the total number of index entries. If the percentage is greater than 20, you need to run REORG: `(((REORGINSERTS+REORGDELETES)*100)/TOTALENTRIES)>20`
# REORGAPPENDINSERT	INTEGER	The number of index entries that have been inserted since the last REORG, REBUILD INDEX or LOAD REPLACE on the index space or partition that have a key value that is greater than the maximum key value in the index or partition. A null value means the number of inserted index entries is unknown. This value can be used with REORGINSERTS to decide when to adjust the PCTFREE specification for the index. For example, if the ratio of REORGAPPENDINSERT to REORGINSERTS is greater than 10 per cent, you might need to run ALTER INDEX to adjust PCTFREE or run REORG more frequently: `((REORGAPPENDINSERT*100)/REORGINSERTS)>10`
# REORGPSEUDODELETES	INTEGER	The number of index entries that have been pseudo-deleted since the last REORG, REBUILD INDEX, or LOAD REPLACE on the index space or partition. A pseudo-delete is a RID entry that has been marked as deleted. A null value means that the number of pseudo-deleted index entries is unknown. This value can be used to determine if a REORG is necessary. For example, if the ratio of pseudo-deletes to total index entries is greater than 10 per cent, you might need to run REORG: `((REORGPSEUDODELETES*100)/TOTALENTRIES)>10`
# REORGMASSDELETE	INTEGER	The number of times that an index or index space partition was mass deleted since the last REORG, REBUILD INDEX, or LOAD REPLACE. A null value means that the number of mass deletes is unknown. If this value is non-zero, a REORG might be necessary.
# REORGLEAFNEAR	INTEGER	The number of index page splits that occurred since the last REORG, REBUILD INDEX, or LOAD REPLACE in which the higher part of the split page was near the location of the original page. The higher part of a split page is near the original page if the two page numbers differ by 16 or less. A null value means that the number of split pages near their original pages is unknown.

Table 143. Descriptions of columns in the INDEXSPACESTATS table (continued)

# Column name	Data type	Description
# REORGLEAFFAR	INTEGER	The number of index page splits that occurred since the last REORG, REBUILD INDEX, or LOAD REPLACE in which the higher part of the split page was far from the location of the original page. The higher part of a split page is far from the original page if the two page numbers differ by more than 16. A null value means that the number of split pages that are far from their original pages is unknown. This value can be used to decide when to run REORG. For example, calculate the ratio of index page splits in which the higher part of the split page was far from the location of the original page to the number of active pages. If this value is greater than 10 per cent, you might need to run REORG: `((REORGLEAFFAR*100)/NACTIVE)>10`
# REORGNUMLEVELS	INTEGER	The number of levels in the index tree that were added or removed since the last REORG, REBUILD INDEX, or LOAD REPLACE. A null value means that the number of added or deleted levels is unknown. If this value has increased since the last REORG, REBUILD INDEX, or LOAD REPLACE, you need to check other values such as REORGPSEUDODELETES to determine whether to run REORG. If this value is less than zero, the index space contains empty pages. Running REORG can save disk space and decrease index sequential scan I/O time by eliminating those empty pages.
# STATSLASTTIME	TIMESTAMP	The timestamp of the last RUNSTATS on the index space or partition. A null value means RUNSTATS has never been run on the index space or partition, or the timestamp of the last RUNSTATS is unknown. You can compare this timestamp to the timestamp of the last REORG on the same object to determine when RUNSTATS is needed. If the date of the last REORG is more recent than the last RUNSTATS, you might need to run RUNSTATS: `(JULIAN_DAY(REORGLASTTIME)>JULIAN_DAY(STATSLASTTIME))`
# STATSINSERTS	INTEGER	The number index entries that have been inserted since the last RUNSTATS on the index space or partition. A null value means that the number of inserted index entries unknown.

Table 143. Descriptions of columns in the INDEXSPACESTATS table (continued)

Column name	Data type	Description
STATSDELETES	INTEGER	The number of index entries that have been deleted since the last RUNSTATS on the index space or partition. A null value means that the number of deleted index entries unknown. This value can be used with STATSINSERTS to determine if RUNSTATS is necessary. For example, suppose that your site's maintenance policies require that RUNSTATS is run after 20 per cent of the rows in an index space have changed. To determine if RUNSTATS is required, calculate the sum of inserted and deleted index entries since the last RUNSTATS. Then calculate the ratio of that sum to the total number of index entries. If the percentage is greater than 20, you need to run RUNSTATS: `(((STATSINSERTS+STATSDELETES)*100)/TOTALENTRIES)>20`
STATSMASSDELETE	INTEGER	The number of times that the index or index space partition was mass deleted since the last RUNSTATS. A null value means that the number of mass deletes is unknown. If this value is non-zero, RUNSTATS might be necessary.
COPYLASTTIME	TIMESTAMP	The timestamp of the last full image copy on the index space or partition. A null value means COPY has never been run on the index space or partition, or the timestamp of the last full image copy is unknown. You can compare this timestamp to the timestamp of the last REORG on the same object to determine when a COPY is needed. If the date of the last REORG is more recent than the last COPY, you might need to run COPY: `(JULIAN_DAY(REORGRLASTTIME)>JULIAN_DAY(COPYLASTTIME))`
COPYUPDATEDPAGES	INTEGER	The number of distinct pages that have been updated since the last COPY. A null value means that the number of updated pages is unknown, or the index was created with COPY NO. You can compare this value to the total number of pages to determine when a COPY is needed. For example, you might want to take a full image copy when 20 percent of the pages have changed: `((COPYUPDATEDPAGES*100)/NACTIVE)>20`
COPYCHANGES	INTEGER	The number of insert delete operations since the last COPY. A null value means that the number of insert or update operations is unknown, or the index was created with COPY NO. This number indicates the approximate number of log records that DB2 processes to recover to the current state. For example, you might want to take a full image copy when DB2 processes more than 10 percent of the index entries from the logs: `((COPYCHANGES*100)/TOTALENTRIES)>10`

Table 143. Descriptions of columns in the INDEXSPACESTATS table (continued)

Column name	Data type	Description
COPYUPDATELRSN	CHAR(6) FOR BIT DATA	The LRSN or RBA of the first update after the last COPY. A null value means that the LRSN or RBA is unknown, or the index was created with COPY NO. Consider running COPY if this value is not in the active logs. To determine the oldest LRSN or RBA in the active logs, use the Print Log Map utility (DSNJU004).
COPYUPDATETIME	TIMESTAMP	The timestamp of the first update after the last COPY. A null value means that the timestamp is unknown, or the index was created with COPY NO. This value has a similar purpose to COPYUPDATELRSN.

Operating with real-time statistics

To use the real-time statistics effectively, you need to understand when DB2 collects and externalizes them, and what factors in your system can affect the statistics. This section contains the following topics:

- "When DB2 externalizes real-time statistics"
- "How DB2 utilities affect the real-time statistics" on page 714
- "How non-DB2 utilities affect real-time statistics" on page 720
- "Real-time statistics on objects in work file databases and the TEMP database" on page 721
- "Real-time statistics on read-only objects" on page 721
- "How dropping objects affects real-time statistics" on page 721
- "How SQL operations affect real-time statistics counters" on page 721
- "Real-time statistics in data sharing" on page 722
- "Improving concurrency with real-time statistics" on page 722
- "Recovering the real-time statistics tables" on page 722
- "Statistics accuracy" on page 722

When DB2 externalizes real-time statistics

DB2 externalizes real-time statistics at the following times:

- When you issue -STOP DATABASE(DSNRTSDB)

  This command stops the in-memory statistics database and externalizes statistics for all objects in the subsystem.
- When you issue -STOP DATABASE(*database-name*) SPACENAM(*space-name*)

  This command externalizes statistics only for *database-name* and *space-name*.
- At the end of the time interval that you specify during installation

  See "Setting the interval for writing real-time statistics" on page 700 for information on how to set this time interval.
- When you issue -STOP DB2 MODE(QUIESCE)

  DB2 writes any statistics that are in memory when you issue this command to the statistics tables. However, if you issue -STOP DB2 MODE(FORCE), DB2 does not write the statistics, and you lose them.
- During utility operations

  "How DB2 utilities affect the real-time statistics" on page 714 gives details on how the utilities modify the statistics tables.

  DB2 does not maintain real-time statistics for any objects in the real-time statistics database. Therefore, if you run a utility with a utility list, and the list

contains any real-time statistics objects, DB2 does not externalize real-time statistics during the execution of that utility for any of the objects in the utility list.

Recommendation: Do not include real-time statistics objects in utility lists.

DB2 does not externalize real-time statistics at a tracker site.

How DB2 utilities affect the real-time statistics

In general, SQL INSERT, UPDATE, and DELETE statements cause DB2 to modify the real-time statistics. However, certain DB2 utilities also affect the statistics. The following section discuss the affect of each of those utilities on the statistics.

How LOAD affects real-time statistics

Table 144 shows how running LOAD on a table space or table space partition affects the TABLESPACESTATS statistics.

Table 144. Changed TABLESPACESTATS values during LOAD

Column name	Settings for LOAD REPLACE after RELOAD phase
TOTALROWS	Number of rows or LOBs loaded[3]
NACTIVE	Actual value
SPACE	Actual value
EXTENTS	Actual value
LOADRLASTTIME	Current timestamp
REORGINSERTS	0
REORGDELETES	0
REORGUPDATES	0
REORGDISORGLOB	0
REORGUNCLUSTINS	0
REORGMASSDELETE	0
REORGNEARINDREF	0
REORGFARINDEF	0
STATSLASTTIME	Current timestamp[1]
STATSINSERTS	0[1]
STATSDELETES	0[1]
STATSUPDATES	0[1]
STATSMASSDELETE	0[1]
COPYLASTTIME	Current timestamp[2]
COPYUPDATEDPAGES	0[2]
COPYCHANGES	0[2]
COPYUPDATELRSN	Null[2]
COPYUPDATETIME	Null[2]

Table 144. Changed TABLESPACESTATS values during LOAD (continued)

Column name	Settings for LOAD REPLACE after RELOAD phase

Note:

1. DB2 sets this value only if the LOAD invocation includes the STATISTICS option.
2. DB2 sets this value only if the LOAD invocation includes the COPYDDN option.
3. Under certain conditions, such as a utility restart, the LOAD utility might not have an accurate count of loaded records. In those cases, DB2 sets this value to null. Some rows that are loaded into a table space and are included in this value might later be removed during the index validation phase or the referential integrity check. DB2 includes counts of those removed records in the statistics that record deleted records.

Table 145 shows how running LOAD affects the INDEXSPACESTATS statistics for an index space or physical index partition.

Table 145. Changed INDEXSPACESTATS values during LOAD

Column name	Settings for LOAD REPLACE after BUILD phase
TOTALENTRIES	Number of index entries added[3]
NLEVELS	Actual value
NACTIVE	Actual value
SPACE	Actual value
EXTENTS	Actual value
LOADRLASTTIME	Current timestamp
REORGINSERTS	0
REORGDELETES	0
REORGAPPENDINSERT	0
REORGPSEUDODELETES	0
REORGMASSDELETE	0
REORGLEAFNEAR	0
REORGLEAFFAR	0
REORGNUMLEVELS	0
STATSLASTTIME	Current timestamp[1]
STATSINSERTS	0[1]
STATSDELETES	0[1]
STATSMASSDELETE	0[1]
COPYLASTTIME	Current timestamp[2]
COPYUPDATEDPAGES	0[2]
COPYCHANGES	0[2]
COPYUPDATELRSN	Null[2]
COPYUPDATETIME	Null[2]

Notes:

1. DB2 sets this value only if the LOAD invocation includes the STATISTICS option.
2. DB2 sets this value only if the LOAD invocation includes the COPYDDN option.
3. Under certain conditions, such as a utility restart, the LOAD utility might not have an accurate count of loaded records. In those cases, DB2 sets this value to null.

For a logical index partition:

- DB2 does not reset the nonpartitioning index when it does a LOAD REPLACE on a partition. Therefore, DB2 does not reset the statistics for the index. The REORG counters from the last REORG are still correct. DB2 updates LOADRLASTTIME when the entire nonpartitioning index is replaced.
- When DB2 does a LOAD RESUME YES on a partition, after the BUILD phase, DB2 increments TOTALENTRIES by the number of index entries that were inserted during the BUILD phase.

How REORG affects real-time statistics

Table 146 shows how running REORG on a table space or table space partition affects the TABLESPACESTATS statistics.

Table 146. Changed TABLESPACESTATS values during REORG

Column name	Settings for REORG SHRLEVEL NONE after RELOAD phase	Settings for REORG SHRLEVEL REFERENCE or CHANGE after SWITCH phase
TOTALROWS	Number rows or LOBs loaded[3]	For SHRLEVEL REFERENCE: Number of rows or LOBs loaded during RELOAD phase For SHRLEVEL CHANGE: Number of rows or LOBs loaded during RELOAD phase +Number of rows inserted during LOG APPLY phase-Number of rows deleted during LOG phase
NACTIVE	Actual value	Actual value
SPACE	Actual value	Actual value
EXTENTS	Actual value	Actual value
REORGLASTTIME	Current timestamp	Current timestamp
REORGINSERTS	0	Actual value[4]
REORGDELETES	0	Actual value[4]
REORGUPDATES	0	Actual value[4]
REORGDISORGLOB	0	Actual value[4]
REORGUNCLUSTINS	0	Actual value[4]
REORGMASSDELETE	0	Actual value[4]
REORGNEARINDREF	0	Actual value[4]
REORGFARINDEF	0	Actual value[4]
STATSLASTTIME	Current timestamp[1]	Current timestamp[1]
STATSINSERTS	0[1]	Actual value[4]
STATSDELETES	0[1]	Actual value[4]
STATSUPDATES	0[1]	Actual value[4]
STATSMASSDELETE	0[1]	Actual value[4]
COPYLASTTIME	Current timestamp[2]	Current timestamp
COPYUPDATEDPAGES	0[2]	Actual value[4]
COPYCHANGES	0[2]	Actual value[4]
COPYUPDATELRSN	Null[2]	Actual value[5]

Table 146. Changed TABLESPACESTATS values during REORG (continued)

Column name	Settings for REORG SHRLEVEL NONE after RELOAD phase	Settings for REORG SHRLEVEL REFERENCE or CHANGE after SWITCH phase
COPYUPDATETIME	Null[2]	Actual value[5]

Notes:

1. DB2 sets this value only if the REORG invocation includes the STATISTICS option.
2. DB2 sets this value only if the REORG invocation includes the COPYDDN option.
3. Under certain conditions, such as a utility restart, the REORG utility might not have an accurate count of loaded records. In those cases, DB2 sets this value to null. Some rows that are loaded into a table space and are included in this value might later be removed during the index validation phase or the referential integrity check. DB2 includes counts of those removed records in the statistics that record deleted records.
4. This is the actual number of inserts, updates, or deletes that are due to applying the log to the shadow copy.
5. This is the LRSN or timestamp for the first update that is due to applying the log to the shadow copy.

Table 147 shows how running REORG affects the INDEXSPACESTATS statistics for an index space or physical index partition.

Table 147. Changed INDEXSPACESTATS values during REORG

Column name	Settings for REORG SHRLEVEL NONE after RELOAD phase	Settings for REORG SHRLEVEL REFERENCE or CHANGE after SWITCH phase
TOTALENTRIES	Number of index entries added[3]	For SHRLEVEL REFERENCE: Number of index entries added during BUILD phase For SHRLEVEL CHANGE: Number of index entries added during BUILD phase +Number of index entries added during LOG phase-Number of index entries deleted during LOG phase
NLEVELS	Actual value	Actual value
NACTIVE	Actual value	Actual value
SPACE	Actual value	Actual value
EXTENTS	Actual value	Actual value
REORGLASTTIME	Current timestamp	Current timestamp
REORGINSERTS	0	Actual value[4]
REORGDELETES	0	Actual value[4]
REORGAPPENDINSERT	0	Actual value[4]
REORGPSEUDODELETES	0	Actual value[4]
REORGMASSDELETE	0	Actual value[4]
REORGLEAFNEAR	0	Actual value[4]
REORGLEAFFAR	0	Actual value[4]
REORGNUMLEVELS	0	Actual value[4]
STATSLASTTIME	Current timestamp[1]	Current timestamp[1]
STATSINSERTS	0[1]	Actual value[4]
STATSDELETES	0[1]	Actual value[4]
STATSMASSDELETE	0[1]	Actual value[4]
COPYLASTTIME	Current timestamp[2]	Unchanged[5]

Table 147. Changed INDEXSPACESTATS values during REORG (continued)

Column name	Settings for REORG SHRLEVEL NONE after RELOAD phase	Settings for REORG SHRLEVEL REFERENCE or CHANGE after SWITCH phase
COPYUPDATEDPAGES	0[2]	Unchanged[5]
COPYCHANGES	0[2]	Unchanged[5]
COPYUPDATELRSN	Null[2]	Unchanged[5]
COPYUPDATETIME	Null[2]	Unchanged[5]

Notes:

1. DB2 sets this value only if the REORG invocation includes the STATISTICS option.
2. DB2 sets this value only if the REORG invocation includes the COPYDDN option.
3. Under certain conditions, such as a utility restart, the REORG utility might not have an accurate count of loaded records. In those cases, DB2 sets this value to null.
4. This is the actual number of inserts, updates, or deletes that are due to applying the log to the shadow copy.
5. Inline COPY is not allowed for SHRLEVEL CHANGE or SHRLEVEL REFERENCE.

For a logical index partition: DB2 does not reset the nonpartitioning index when it does a REORG on a partition. Therefore, DB2 does not reset the statistics for the index. The REORG counters from the last REORG are still correct. DB2 updates REORGLASTTIME when the entire nonpartitioning index is reorganized.

How REBUILD INDEX affects real-time statistics

Table 148 shows how running REBUILD INDEX affects the INDEXSPACESTATS statistics for an index space or physical index partition.

Table 148. Changed INDEXSPACESTATS values during REBUILD INDEX

Column name	Settings after BUILD phase
TOTALENTRIES	Number of index entries added[1]
NLEVELS	Actual value
NACTIVE	Actual value
SPACE	Actual value
EXTENTS	Actual value
REBUILDLASTTIME	Current timestamp
REORGINSERTS	0
REORGDELETES	0
REORGAPPENDINSERT	0
REORGPSEUDODELETES	0
REORGMASSDELETE	0
REORGLEAFNEAR	0
REORGLEAFFAR	0
REORGNUMLEVELS	0

Note:

1. Under certain conditions, such as a utility restart, the REBUILD utility might not have an accurate count of loaded records. In those cases, DB2 sets this value to null.

For a logical index partition: DB2 does not collect TOTALENTRIES statistics for the entire nonpartitioning index when it runs REBUILD INDEX. Therefore, DB2 does not

reset the statistics for the index. The REORG counters from the last REORG are still correct. DB2 updates REBUILDLASTTIME when the entire nonpartitioning index is rebuilt.

How RUNSTATS affects real-time statistics

Only RUNSTATS UPDATE ALL affects the real-time statistics. When the RUNSTATS job starts, DB2 externalizes all in-memory statistics to the real-time statistics tables.

Table 149 shows how running RUNSTATS UPDATE ALL on a table space or table space partition affects the TABLESPACESTATS statistics.

Table 149. Changed TABLESPACESTATS values during RUNSTATS UPDATE ALL

Column name	During UTILINIT phase	After RUNSTATS phase
STATSLASTTIME	Current timestamp[1]	Timestamp of the start of RUNSTATS phase
STATSINSERTS	Actual value[1]	Actual value[2]
STATSDELETES	Actual value[1]	Actual value[2]
STATSUPDATES	Actual value[1]	Actual value[2]
STATSMASSDELETE	Actual value[1]	Actual value[2]

Notes:

1. DB2 externalizes the current in-memory values.
2. This value is 0 for SHRLEVEL REFERENCE, or the actual value for SHRLEVEL CHANGE.

Table 150 shows how running RUNSTATS UPDATE ALL on an index affects the INDEXSPACESTATS statistics.

Table 150. Changed INDEXSPACESTATS values during RUNSTATS UPDATE ALL

Column name	During UTILINIT phase	After RUNSTATS phase
STATSLASTTIME	Current timestamp[1]	Timestamp of the start of RUNSTATS phase
STATSINSERTS	Actual value[1]	Actual value[2]
STATSDELETES	Actual value[1]	Actual value[2]
STATSMASSDELETE	Actual value[1]	Actual value[2]

Notes:

1. DB2 externalizes the current in-memory values.
2. This value is 0 for SHRLEVEL REFERENCE, or the actual value for SHRLEVEL CHANGE.

How COPY affects real-time statistics

When a COPY job starts, DB2 externalizes all in-memory statistics to the real-time statistics tables. Statistics are gathered only for a full image copy or an incremental copy, but not for a data set copy.

Table 151 on page 720 shows how running COPY on a table space or table space partition affects the TABLESPACESTATS statistics.

Table 151. Changed TABLESPACESTATS values during COPY

Column name	During UTILINIT phase	After COPY phase
COPYLASTTIME	Current timestamp[1]	Timestamp of the start of COPY phase
COPYUPDATEDPAGES	Actual value[1]	Actual value[2]
COPYCHANGES	Actual value[1]	Actual value[2]
COPYUPDATELRSN	Actual value[1]	Actual value[3]
COPYUPDATETIME	Actual value[1]	Actual value[3]

Notes:

1. DB2 externalizes the current in-memory values.
2. This value is 0 for SHRLEVEL REFERENCE, or the actual value for SHRLEVEL CHANGE.
3. This value is null for SHRLEVEL REFERENCE, or the actual value for SHRLEVEL CHANGE.

Table 152 shows how running COPY on an index affects the INDEXSPACESTATS statistics.

Table 152. Changed INDEXSPACESTATS values during COPY

Column name	During UTILINIT phase	After COPY phase
COPYLASTTIME	Current timestamp[1]	Timestamp of the start of COPY phase
COPYUPDATEDPAGES	Actual value[1]	Actual value[2]
COPYCHANGES	Actual value[1]	Actual value[2]
COPYUPDATELRSN	Actual value[1]	Actual value[3]
COPYUPDATETIME	Actual value[1]	Actual value[3]

Note:

1. DB2 externalizes the current in-memory values.
2. This value is 0 for SHRLEVEL REFERENCE, or the actual value for SHRLEVEL CHANGE.
3. This value is null for SHRLEVEL REFERENCE, or the actual value for SHRLEVEL CHANGE.

How RECOVER affects real-time statistics

After recovery to the current state, the in-memory counter fields are still valid, so DB2 does not modify them. However, after a point-in-time recovery, the statistics might not be valid. DB2 therefore sets all the REORG, STATS, and COPY counter statistics to null after a point-in-time recovery. After recovery to the current state, DB2 sets NACTIVE, SPACE, and EXTENTS to their new values. After a point-in-time recovery, DB2 sets NLEVELS, NACTIVE, SPACE, and EXTENTS to their new values.

How non-DB2 utilities affect real-time statistics

Non-DB2 utilities do not affect real-time statistics. Therefore, an object that is the target of a non-DB2 COPY, LOAD, REBUILD, REORG or RUNSTATS can cause incorrect statistics to be inserted in the real-time statistics tables. Follow this process to ensure correct statistics when you run non-DB2 utilities:

1. Stop the table space or index on which you plan to run the utility. This action causes DB2 to write the in-memory statistics to the real-time statistics tables and initialize the in-memory counters.
2. Run the utility.
3. When the utility completes, update the statistics tables with new totals, timestamps and zero incremental counter values.

Real-time statistics on objects in work file databases and the TEMP database

Although you cannot run utilities on objects in the work files databases and TEMP database, DB2 records the NACTIVE, SPACE, and EXTENTS statistics on table spaces in those databases.

Real-time statistics on read-only objects

DB2 does not externalize the NACTIVE, SPACE, or EXTENTS statistics for read-only objects.

How dropping objects affects real-time statistics

If you drop a table space or index, DB2 deletes its statistics from the real-time statistics tables. However, if the real-time statistics database is not available when you drop a table space or index, the statistics remain in the real-time statistics tables, even though the corresponding object no longer exists. You need to use SQL DELETE statements to manually remove those rows from the real-time statistics tables.

If a row still exists in the real-time statistics tables for a dropped table space or index, and if you create a new object with the same DBID and PSID as the dropped object, DB2 reinitializes the row before it updates any values in that row.

How SQL operations affect real-time statistics counters

SQL operations affect the counter columns in the real-time statistics tables. These are the columns that record that record the number of insert, delete, or update operations, as well as total counters TOTALROWS and TOTALENTRIES.

UPDATE: When you perform an UPDATE, DB2 increments the update counters.

INSERT: When you perform an INSERT, DB2 increments the insert counters. DB2 keeps separate counters for clustered and unclustered UPDATEs.

DELETE: When you perform a DELETE, DB2 increments the delete counters.

ROLLBACK: When you perform a ROLLBACK, DB2 increments the the counters, depending on the type of SQL operation that is rolled back:

Rolled-back SQL statement	Incremented counters
UPDATE	Update counters
INSERT	Delete counters
DELETE	Insert counters

Mass DELETE: When you perform a mass delete operation on a table space does not cause DB2 to reset the counter columns in the real-time statistics tables. After a

mass delete operation, the value in a counter column includes the count from before the mass delete operation, as well as the count after the mass delete operation.

Real-time statistics in data sharing

In a data sharing environment, DB2 members update their statistics serially. Each member reads the target row from the statistics table, obtains a lock, aggregates its in-memory statistics, and updates the statistics table with the new totals. Each member sets its own interval for writing real-time statistics.

DB2 does locking based on the lock size of the DSNRTSDB.DSNRTSTS table space. DB2 uses cursor stability isolation and CURRENTDATA(YES) when it reads the statistics tables.

At the beginning of a RUNSTATS job, all data sharing members externalize their statistics to the real-time statistics tables and reset their in-memory statistics. If all members cannot externalize their statistics, DB2 sets STATSLASTTIME to null. An error in gathering and externalizing statistics does not prevent RUNSTATS from running.

At the beginning of a COPY job, all data sharing members externalize their statistics to the real-time statistics tables and reset their in-memory statistics. If all members cannot externalize their statistics, DB2 sets COPYLASTTIME to null. An error in gathering and externalizing statistics does not prevent COPY from running.

Utilities that reset page sets to empty can invalidate the in-memory statistics of other DB2 members. The member that resets a page set notifies the other DB2 members that a page set has been reset to empty, and the in-memory stats are invalidated. If the notify process fails, the utility that resets the page set does not fail. DB2 sets the appropriate timestamp (REORGLASTTIME, STATSLASTTIME or COPYLASTTIME) to null in the row for the empty page set to indicate that the statistics for that page set are unknown.

Improving concurrency with real-time statistics

Follow these recommendations to reduce the risk of timeouts and deadlocks when you work with the real-time statistics tables:

- When you run COPY, RUNSTATS or REORG on the real-time statistics objects, use SHRLEVEL CHANGE.
- When you execute SQL statements to query the real-time statistics tables, use uncommitted read isolation.

Recovering the real-time statistics tables

When you recover a DB2 subsystem after a disaster, you need to perform the following actions on the real-time statistics database:

- Recover the real-time statistics objects after you recover the DB2 catalog and directory.
- Start the real-time statistics database explicitly, after DB2 restart.

Statistics accuracy

In general, the real-time statistics are accurate values. However, several factors can affect the accuracy of the statistics:

- Certain utility restart scenarios
- A DB2 subsystem failure

- A notify failure in a data sharing environment

If you think that some statistics values might be inaccurate, you can correct the statistics by running REORG, RUNSTATS, or COPY on the objects for which DB2 generated the statistics.

Notices

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

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

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

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

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

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

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

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs

© Copyright IBM Corp. 1983, 2001

and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J74/G4
555 Bailey Avenue
P.O. Box 49023
San Jose, CA 95161-9023
U.S.A.

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

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

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

COPYRIGHT LICENSE:

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

Programming interface information

This book is intended to help you to use DB2 for OS/390 and z/OS utilities.

This book also documents General-use Programming Interface and Associated Guidance Information and Product-sensitive Programming Interface and Associated Guidance Information provided by IBM DATABASE 2 Universal Database Server for OS/390 and z/OS (DB2 for OS/390 and z/OS).

General-use programming interfaces allow the customer to write programs that obtain the services of DB2 for OS/390 and z/OS.

General-use Programming Interface and Associated Guidance Information is identified where it occurs, by an entry in a column of a table.

Product-sensitive programming interfaces allow the customer installation to perform tasks such as diagnosing, modifying, monitoring, repairing, tailoring, or tuning of this IBM software product. Use of such interfaces creates dependencies on the detailed design or implementation of the IBM software product. Product-sensitive programming interfaces should be used only for these specialized purposes. Because of their dependencies on detailed design and implementation, it is to be

expected that programs written to such interfaces may require changes in order to run with new product releases or versions, or as a result of service.

Product-sensitive Programming Interface and Associated Guidance Information is identified where it occurs by an entry in a column of a table, or by the following marking:

```
Product-sensitive Programming Interface

Product-sensitive Programming Interface and Associated Guidance Information ...

End of Product-sensitive Programming Interface
```

Trademarks

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

3090	ES/3090
AS/400	IBM
AD/Cycle	IBM Registry
APL2	IMS
AS/400	Language Environment
BookManager	MVS/DFP
C/370	MVS/ESA
CICS	Net.Data
CICS/ESA	OpenEdition
CICS/ESA	OS/2
DATABASE 2	OS/390
DataHub	OS/400
DataPropagator	Parallel Sysplex
DB2	PR/SM
DB2 Connect	QMF
DB2 Universal Database	RACF
DFSMS/MVS	RAMAC
DFSMSdfp	RETAIN
DFSMSdss	RMF
DFSMShsm	S/390
DFSORT	SAA
Distributed Relational Database Architecture	SecureWay
DRDA	SQL/DS
DXT	System/370
Enterprise Storage Server	System/390
Enterprise System/3090	VTAM
Enterprise System/9000	WebSphere

JDBC™ and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft® is a registered trademark of Microsoft Corporation in the United States, other countries, or both.

UNIX® is a registered trademark of The Open Group in the United States and other countries.

MCS is a trademark of Intel Corporation in the United States, other countries, or both.

SET™ is a trademark owned by SET Secure Electronic Transaction LLC.

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

Glossary

The following terms and abbreviations are defined as they are used in the DB2 library.

A

abend. Abnormal end of task.

abend reason code. A 4-byte hexadecimal code that uniquely identifies a problem with DB2. A complete list of DB2 abend reason codes and their explanations is contained in *DB2 Messages and Codes*.

abnormal end of task (abend). Termination of a task, job, or subsystem because of an error condition that recovery facilities cannot resolve during execution.

access method services. The facility that is used to define and reproduce VSAM key-sequenced data sets.

access path. The path that is used to locate data that is specified in SQL statements. An access path can be indexed or sequential.

active log. The portion of the DB2 log to which log records are written as they are generated. The active log always contains the most recent log records, whereas the archive log holds those records that are older and no longer fit on the active log.

address space. A range of virtual storage pages that is identified by a number (ASID) and a collection of segment and page tables that map the virtual pages to real pages of the computer's memory.

address space connection. The result of connecting an allied address space to DB2. Each address space that contains a task that is connected to DB2 has exactly one address space connection, even though more than one task control block (TCB) can be present. See also *allied address space* and *task control block*.

agent. As used in DB2, the structure that associates all processes that are involved in a DB2 unit of work. An *allied agent* is generally synonymous with an *allied thread*. *System agents* are units of work that process independently of the allied agent, such as prefetch processing, deferred writes, and service tasks.

alias. An alternative name that can be used in SQL statements to refer to a table or view in the same or a remote DB2 subsystem.

allied address space. An area of storage that is external to DB2 and that is connected to DB2. An allied address space is capable of requesting DB2 services.

allied thread. A thread that originates at the local DB2 subsystem and that can access data at a remote DB2 subsystem.

ambiguous cursor. A database cursor that is not defined with the FOR FETCH ONLY clause or the FOR UPDATE OF clause, is not defined on a read-only result table, is not the target of a WHERE CURRENT clause on an SQL UPDATE or DELETE statement, and is in a plan or package that contains either PREPARE or EXECUTE IMMEDIATE SQL statements.

American National Standards Institute (ANSI). An organization consisting of producers, consumers, and general interest groups, that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States.

ANSI. American National Standards Institute.

API. Application programming interface.

APPL. A VTAM network definition statement that is used to define DB2 to VTAM as an application program that uses SNA LU 6.2 protocols.

application. A program or set of programs that performs a task; for example, a payroll application.

application plan. The control structure that is produced during the bind process. DB2 uses the application plan to process SQL statements that it encounters during statement execution.

application process. The unit to which resources and locks are allocated. An application process involves the execution of one or more programs.

application programming interface (API). A functional interface that is supplied by the operating system or by a separately orderable licensed program that allows an application program that is written in a high-level language to use specific data or functions of the operating system or licensed program.

application server. The target of a request from a remote application. In the DB2 environment, the application server function is provided by the distributed data facility and is used to access DB2 data from remote applications.

archive log. The portion of the DB2 log that contains log records that have been copied from the active log.

ASCII. An encoding scheme that is used to represent strings in many environments, typically on PCs and workstations. Contrast with *EBCDIC* and *Unicode*.

attachment facility. An interface between DB2 and TSO, IMS, CICS, or batch address spaces. An attachment facility allows application programs to access DB2.

© Copyright IBM Corp. 1983, 2001

attribute. A characteristic of an entity. For example, in database design, the phone number of an employee is one of that employee's attributes.

authorization ID. A string that can be verified for connection to DB2 and to which a set of privileges is allowed. It can represent an individual, an organizational group, or a function, but DB2 does not determine this representation.

auxiliary index. An index on an auxiliary table in which each index entry refers to a LOB.

auxiliary table. A table that stores columns outside the table in which they are defined. Contrast with *base table*.

B

backward log recovery. The fourth and final phase of restart processing during which DB2 scans the log in a backward direction to apply UNDO log records for all aborted changes.

base table. (1) A table that is created by the SQL CREATE TABLE statement and that holds persistent data. Contrast with *result table* and *temporary table*.

(2) A table containing a LOB column definition. The actual LOB column data is not stored with the base table. The base table contains a row identifier for each row and an indicator column for each of its LOB columns. Contrast with *auxiliary table*.

base table space. A table space that contains base tables.

basic sequential access method (BSAM). An access method for storing or retrieving data blocks in a continuous sequence, using either a sequential access or a direct access device.

binary integer. A basic data type that can be further classified as small integer or large integer.

binary large object (BLOB). A sequence of bytes where the size of the value ranges from 0 bytes to 2 GB–1. Such a string does not have an associated CCSID.

bind. The process by which the output from the SQL precompiler is converted to a usable control structure, often called an access plan, application plan, or package. During this process, access paths to the data are selected and some authorization checking is performed. The types of bind are:

automatic bind. (More correctly, *automatic rebind*) A process by which SQL statements are bound automatically (without a user issuing a BIND command) when an application process begins execution and the bound application plan or package it requires is not valid.

dynamic bind. A process by which SQL statements are bound as they are entered.

incremental bind. A process by which SQL statements are bound during the execution of an application process, because they could not be bound during the bind process, and VALIDATE(RUN) was specified.

static bind. A process by which SQL statements are bound after they have been precompiled. All static SQL statements are prepared for execution at the same time.

BLOB. Binary large object.

BMP. Batch Message Processing (IMS).

bootstrap data set (BSDS). A VSAM data set that contains name and status information for DB2, as well as RBA range specifications, for all active and archive log data sets. It also contains passwords for the DB2 directory and catalog, and lists of conditional restart and checkpoint records.

BSAM. Basic sequential access method.

BSDS. Bootstrap data set.

buffer pool. Main storage that is reserved to satisfy the buffering requirements for one or more table spaces or indexes.

built-in function. A function that DB2 supplies. Contrast with *user-defined function*.

C

CAF. Call attachment facility.

call attachment facility (CAF). A DB2 attachment facility for application programs that run in TSO or MVS batch. The CAF is an alternative to the DSN command processor and provides greater control over the execution environment.

cascade delete. The way in which DB2 enforces referential constraints when it deletes all descendent rows of a deleted parent row.

cast function. A function that is used to convert instances of a (source) data type into instances of a different (target) data type. In general, a cast function has the name of the target data type. It has one single argument whose type is the source data type; its return type is the target data type.

catalog. In DB2, a collection of tables that contains descriptions of objects such as tables, views, and indexes.

catalog table. Any table in the DB2 catalog.

CCSID. Coded character set identifier.

CDB. Communications database.

character large object (CLOB). A sequence of bytes representing single-byte characters or a mixture of single- and double-byte characters where the size of the value can be up to 2 GB–1. In general, character large object values are used whenever a character string might exceed the limits of the VARCHAR type.

character set. A defined set of characters.

character string. A sequence of bytes that represent bit data, single-byte characters, or a mixture of single-byte and multibyte characters.

check constraint. See *table check constraint*.

check integrity. The condition that exists when each row in a table conforms to the table check constraints that are defined on that table. Maintaining check integrity requires DB2 to enforce table check constraints on operations that add or change data.

check pending. A state of a table space or partition that prevents its use by some utilities and some SQL statements because of rows that violate referential constraints, table check constraints, or both.

checkpoint. A point at which DB2 records internal status information on the DB2 log; the recovery process uses this information if DB2 abnormally terminates.

CI. Control interval.

CICS. Represents (in this publication) one of the following products:

CICS Transaction Server for OS/390: Customer Information Control System Transaction Server for OS/390
CICS/ESA: Customer Information Control System/Enterprise Systems Architecture
CICS/MVS: Customer Information Control System/Multiple Virtual Storage

CICS attachment facility. A DB2 subcomponent that uses the MVS subsystem interface (SSI) and cross storage linkage to process requests from CICS to DB2 and to coordinate resource commitment.

CIDF. Control interval definition field.

claim. A notification to DB2 that an object is being accessed. Claims prevent drains from occurring until the claim is released, which usually occurs at a commit point. Contrast with *drain*.

claim class. A specific type of object access that can be one of the following:

Cursor stability (CS)
Repeatable read (RR)
Write

claim count. A count of the number of agents that are accessing an object.

clause. In SQL, a distinct part of a statement, such as a SELECT clause or a WHERE clause.

CLIST. Command list. A language for performing TSO tasks.

CLOB. Character large object.

clustering index. An index that determines how rows are physically ordered in a table space.

coded character set. A set of unambiguous rules that establish a character set and the one-to-one relationships between the characters of the set and their coded representations.

coded character set identifier (CCSID). A 16-bit number that uniquely identifies a coded representation of graphic characters. It designates an encoding scheme identifier and one or more pairs consisting of a character set identifier and an associated code page identifier.

cold start. A process by which DB2 restarts without processing any log records. Contrast with *warm start*.

collection. A group of packages that have the same qualifier.

column. The vertical component of a table. A column has a name and a particular data type (for example, character, decimal, or integer).

column function. An operation that derives its result by using values from one or more rows. Contrast with *scalar function*.

command. A DB2 operator command or a DSN subcommand. A command is distinct from an SQL statement.

command recognition character (CRC). A character that permits an MVS console operator or an IMS subsystem user to route DB2 commands to specific DB2 subsystems.

commit. The operation that ends a unit of work by releasing locks so that the database changes that are made by that unit of work can be perceived by other processes.

commit point. A point in time when data is considered consistent.

committed phase. The second phase of the multisite update process that requests all participants to commit the effects of the logical unit of work.

communications database (CDB). A set of tables in the DB2 catalog that are used to establish conversations with remote database management systems.

compression dictionary. The dictionary that controls the process of compression and decompression. This dictionary is created from the data in the table space or table space partition.

concurrency. The shared use of resources by more than one application process at the same time.

conditional restart. A DB2 restart that is directed by a user-defined conditional restart control record (CRCR).

connection. In SNA, the existence of a communication path between two partner LUs that allows information to be exchanged (for example, two DB2 subsystems that are connected and communicating by way of a conversation).

connection ID. An identifier that is supplied by the attachment facility and that is associated with a specific address space connection.

consistency token. A timestamp that is used to generate the version identifier for an application. See also *version*.

constant. A language element that specifies an unchanging value. Constants are classified as string constants or numeric constants. Contrast with *variable*.

constraint. A rule that limits the values that can be inserted, deleted, or updated in a table. See *referential constraint*, *table check constraint*, and *uniqueness constraint*.

control interval (CI). A fixed-length area or direct access storage in which VSAM stores records and creates distributed free space. Also, in a key-sequenced data set or file, the set of records pointed to by an entry in the sequence-set index record. The control interval is the unit of information that VSAM transmits to or from direct access storage. A control interval always includes an integral number of physical records.

control interval definition field (CIDF). In VSAM, a field located in the 4 bytes at the end of each control interval; it describes the free space, if any, in the control interval.

conversation. Communication, which is based on LU 6.2 or Advanced Program-to-Program Communication (APPC), between an application and a remote transaction program over an SNA logical unit-to-logical unit (LU-LU) session that allows communication while processing a transaction.

coordinator. The system component that coordinates the commit or rollback of a unit of work that includes work that is done on one or more other systems.

correlation ID. An identifier that is associated with a specific thread. In TSO, it is either an authorization ID or the job name.

CRC. Command recognition character.

CRCR. Conditional restart control record. See also *conditional restart*.

created temporary table. A table that holds temporary data and is defined with the SQL statement CREATE GLOBAL TEMPORARY TABLE. Information about created temporary tables is stored in the DB2 catalog, so this kind of table is persistent and can be shared across application processes. Contrast with *declared temporary table*. See also *temporary table*.

CT. Cursor table.

current data. Data within a host structure that is current with (identical to) the data within the base table.

current status rebuild. The second phase of restart processing during which the status of the subsystem is reconstructed from information on the log.

cursor table (CT). The copy of the skeleton cursor table that is used by an executing application process.

cycle. A set of tables that can be ordered so that each table is a descendent of the one before it, and the first table is a descendent of the last table. A self-referencing table is a cycle with a single member.

D

DASD. Direct access storage device.

database. A collection of tables, or a collection of table spaces and index spaces.

database access thread. A thread that accesses data at the local subsystem on behalf of a remote subsystem.

database administrator (DBA). An individual who is responsible for designing, developing, operating, safeguarding, maintaining, and using a database.

database descriptor (DBD). An internal representation of a DB2 database definition, which reflects the data definition that is in the DB2 catalog. The objects that are defined in a database descriptor are table spaces, tables, indexes, index spaces, and relationships.

database management system (DBMS). A software system that controls the creation, organization, and modification of a database and the access to the data stored within it.

database request module (DBRM). A data set member that is created by the DB2 precompiler and that contains information about SQL statements. DBRMs are used in the bind process.

database server. The target of a request from a local application or an intermediate database server. In the DB2 environment, the database server function is provided by the distributed data facility to access DB2 data from local applications, or from a remote database server that acts as an intermediate database server.

DATABASE 2 Interactive (DB2I). The DB2 facility that provides for the execution of SQL statements, DB2 (operator) commands, programmer commands, and utility invocation.

data currency. The state in which data that is retrieved into a host variable in your program is a copy of data in the base table.

data definition name (ddname). The name of a data definition (DD) statement that corresponds to a data control block containing the same name.

Data Language/I (DL/I). The IMS data manipulation language; a common high-level interface between a user application and IMS.

data partition. A VSAM data set that is contained within a partitioned table space.

data type. An attribute of columns, literals, host variables, special registers, and the results of functions and expressions.

date. A three-part value that designates a day, month, and year.

date duration. A decimal integer that represents a number of years, months, and days.

DBA. Database administrator.

DBCLOB. Double-byte character large object.

DBCS. Double-byte character set.

DBD. Database descriptor.

DBID. Database identifier.

DBMS. Database management system.

DBRM. Database request module.

DB2 catalog. Tables that are maintained by DB2 and contain descriptions of DB2 objects, such as tables, views, and indexes.

DB2 command. An instruction to the DB2 subsystem allowing a user to start or stop DB2, to display information on current users, to start or stop databases, to display information on the status of databases, and so on.

DB2 for VSE & VM. The IBM DB2 relational database management system for the VSE and VM operating systems.

DB2I. DATABASE 2 Interactive.

DCLGEN. Declarations generator.

DDF. Distributed data facility.

ddname. Data definition name.

deadlock. Unresolvable contention for the use of a resource such as a table or an index.

declarations generator (DCLGEN). A subcomponent of DB2 that generates SQL table declarations and COBOL, C, or PL/I data structure declarations that conform to the table. The declarations are generated from DB2 system catalog information. DCLGEN is also a DSN subcommand.

declared temporary table. A table that holds temporary data and is defined with the SQL statement DECLARE GLOBAL TEMPORARY TABLE. Information about declared temporary tables is not stored in the DB2 catalog, so this kind of table is not persistent and can only be used by the application process that issued the DECLARE statement. Contrast with *created temporary table*. See also *temporary table*.

default value. A predetermined value, attribute, or option that is assumed when no other is explicitly specified.

deferred write. The process of asynchronously writing changed data pages to disk.

delimited identifier. A sequence of characters that are enclosed within double quotation marks ("). The sequence must consist of a letter followed by zero or more characters, each of which is a letter, digit, or the underscore character (_).

dependent. An object (row, table, or table space) that has at least one parent. The object is also said to be a dependent (row, table, or table space) of its parent. See *parent row*, *parent table*, *parent table space*.

dependent row. A row that contains a foreign key that matches the value of a primary key in the parent row.

dependent table. A table that is a dependent in at least one referential constraint.

descendent. An object that is a dependent of an object or is the dependent of a descendent of an object.

descendent row. A row that is dependent on another row, or a row that is a descendent of a dependent row.

descendent table. A table that is a dependent of another table, or a table that is a descendent of a dependent table.

direct access storage device (DASD). A device in which access time is independent of the location of the data.

directory. The DB2 system database that contains internal objects such as database descriptors and skeleton cursor tables.

distinct type. A user-defined data type that is internally represented as an existing type (its source type), but is considered to be a separate and incompatible type for semantic purposes.

distributed data facility (DDF). A set of DB2 components through which DB2 communicates with another RDBMS.

Distributed Relational Database Architecture (DRDA). A connection protocol for distributed relational database processing that is used by IBM's relational database products. DRDA includes protocols for communication between an application and a remote relational database management system, and for communication between relational database management systems.

DL/I. Data Language/I.

double-byte character large object (DBCLOB). A sequence of bytes representing double-byte characters where the size of the values can be up to 2 GB. In general, double-byte character large object values are used whenever a double-byte character string might exceed the limits of the VARGRAPHIC type.

double-byte character set (DBCS). A set of characters, which are used by national languages such as Japanese and Chinese, that have more symbols than can be represented by a single byte. Each character is 2 bytes in length. Contrast with *single-byte character set* and *multibyte character set*.

double-precision floating point number. A 64-bit approximate representation of a real number.

drain. The act of acquiring a locked resource by quiescing access to that object.

drain lock. A lock on a claim class that prevents a claim from occurring.

DRDA. Distributed Relational Database Architecture.

DRDA access. An open method of accessing distributed data that you can use to can connect to another database server to execute packages that were previously bound at the server location. You use the SQL CONNECT statement or an SQL statement with a three-part name to identify the server. Contrast with *private protocol access*.

DSN. (1) The default DB2 subsystem name. (2) The name of the TSO command processor of DB2. (3) The first three characters of DB2 module and macro names.

duration. A number that represents an interval of time. See *date duration*, *labeled duration*, and *time duration*.

dynamic SQL. SQL statements that are prepared and executed within an application program while the program is executing. In dynamic SQL, the SQL source is contained in host language variables rather than being coded into the application program. The SQL statement can change several times during the application program's execution.

E

EBCDIC. Extended binary coded decimal interchange code. An encoding scheme that is used to represent character data in the OS/390, MVS, VM, VSE, and OS/400® environments. Contrast with *ASCII* and *Unicode*.

EDM pool. A pool of main storage that is used for database descriptors, application plans, authorization cache, application packages, and dynamic statement caching.

embedded SQL. SQL statements that are coded within an application program. See *static SQL*.

escape character. The symbol that is used to enclose an SQL delimited identifier. The escape character is the double quotation mark ("), except in COBOL applications, where the user assigns the symbol, which is either a double quotation mark or an apostrophe (').

ESDS. Entry sequenced data set.

EUR. IBM European Standards.

exception table. A table that holds rows that violate referential constraints or table check constraints that the CHECK DATA utility finds.

exclusive lock. A lock that prevents concurrently executing application processes from reading or changing data. Contrast with *share lock*.

executable statement. An SQL statement that can be embedded in an application program, dynamically prepared and executed, or issued interactively.

exit routine. A user-written (or IBM-provided default) program that receives control from DB2 to perform specific functions. Exit routines run as extensions of DB2.

external function. A function for which the body is written in a programming language that takes scalar argument values and produces a scalar result for each invocation. Contrast with *sourced function*, *built-in function*, and *SQL function*.

F

fallback. The process of returning to a previous release of DB2 after attempting or completing migration to a current release.

field procedure. A user-written exit routine that is designed to receive a single value and transform (encode or decode) it in any way the user can specify.

fixed-length string. A character or graphic string whose length is specified and cannot be changed. Contrast with *varying-length string*.

foreign key. A column or set of columns in a dependent table of a constraint relationship. The key must have the same number of columns, with the same descriptions, as the primary key of the parent table. Each foreign key value must either match a parent key value in the related parent table or be null.

forward log recovery. The third phase of restart processing during which DB2 processes the log in a forward direction to apply all REDO log records.

free space. The total amount of unused space in a page; that is, the space that is not used to store records or control information is free space.

function. A mapping, embodied as a program (the function body), invocable by means of zero or more input values (arguments), to a single value (the result). See also *column function* and *scalar function*.

Functions can be user-defined, built-in, or generated by DB2. (See *built-in function*, *cast function*, *external function*, *sourced function*, *SQL function*, and *user-defined function*.)

G

GB. Gigabyte (1 073 741 824 bytes).

GBP. Group buffer pool.

generalized trace facility (GTF). An MVS service program that records significant system events such as I/O interrupts, SVC interrupts, program interrupts, or external interrupts.

getpage. An operation in which DB2 accesses a data page.

global lock contention. Conflicts on locking requests between different DB2 members of a data sharing group when those members are trying to serialize shared resources.

governor. See *resource limit facility*.

gross lock. The *shared*, *update*, or *exclusive* mode locks on a table, partition, or table space.

group buffer pool (GBP). A coupling facility cache structure that is used by a data sharing group to cache data and to ensure that the data is consistent for all members.

GTF. Generalized trace facility.

H

help panel. A screen of information presenting tutorial text to assist a user at the terminal.

host language. A programming language in which you can embed SQL statements.

host program. An application program that is written in a host language and that contains embedded SQL statements.

HSM. Hierarchical storage manager.

I

IDCAMS. An IBM program that is used to process access method services commands. It can be invoked as a job or jobstep, from a TSO terminal, or from within a user's application program.

identify. A request that an attachment service program in an address space that is separate from DB2 issues via the MVS subsystem interface to inform DB2 of its existence and to initiate the process of becoming connected to DB2.

identity column. A column that provides a way for DB2 to automatically generate a numeric value for each row. The generated values are unique if cycling is not used. Identity columns are defined with the AS IDENTITY clause. Uniqueness of values can be ensured by defining a single-column unique index using the identity column. A table can have no more than one identity column.

IFCID. Instrumentation facility component identifier.

IFI. Instrumentation facility interface.

IFI call. An invocation of the instrumentation facility interface (IFI) by means of one of its defined functions.

image copy. An exact reproduction of all or part of a table space. DB2 provides utility programs to make full image copies (to copy the entire table space) or incremental image copies (to copy only those pages that have been modified since the last image copy).

IMS. Information Management System.

IMS attachment facility. A DB2 subcomponent that uses MVS subsystem interface (SSI) protocols and cross-memory linkage to process requests from IMS to DB2 and to coordinate resource commitment.

in-abort. A status of a unit of recovery. If DB2 fails after a unit of recovery begins to be rolled back, but before the process is completed, DB2 continues to back out the changes during restart.

in-commit. A status of a unit of recovery. If DB2 fails after beginning its phase 2 commit processing, it "knows," when restarted, that changes made to data are consistent. Such units of recovery are termed *in-commit*.

index. A set of pointers that are logically ordered by the values of a key. Indexes can provide faster access to data and can enforce uniqueness on the rows in a table.

index key. The set of columns in a table that is used to determine the order of index entries.

index partition. A VSAM data set that is contained within a partitioning index space.

index space. A page set that is used to store the entries of one index.

indoubt. A status of a unit of recovery. If DB2 fails after it has finished its phase 1 commit processing and before it has started phase 2, only the commit coordinator knows if an individual unit of recovery is to be committed or rolled back. At emergency restart, if DB2 lacks the information it needs to make this decision, the status of the unit of recovery is *indoubt* until DB2 obtains this information from the coordinator. More than one unit of recovery can be indoubt at restart.

indoubt resolution. The process of resolving the status of an indoubt logical unit of work to either the committed or the rollback state.

inflight. A status of a unit of recovery. If DB2 fails before its unit of recovery completes phase 1 of the commit process, it merely backs out the updates of its unit of recovery at restart. These units of recovery are termed *inflight*.

inline copy. A copy that is produced by the LOAD or REORG utility. The data set that the inline copy produces is logically equivalent to a full image copy that is produced by running the COPY utility with read-only access (SHRLEVEL REFERENCE).

instrumentation facility component identifier (IFCID). A value that names and identifies a trace record of an event that can be traced. As a parameter on the START TRACE and MODIFY TRACE commands, it specifies that the corresponding event is to be traced.

instrumentation facility interface (IFI). A programming interface that enables programs to obtain online trace data about DB2, to submit DB2 commands, and to pass data to DB2.

Interactive System Productivity Facility (ISPF). An IBM licensed program that provides interactive dialog services.

intermediate database server. The target of a request from a local application or a remote application requester that is forwarded to another database server. In the DB2 environment, the remote request is forwarded transparently to another database server if the object that is referenced by a three-part name does not reference the local location.

internal resource lock manager (IRLM). An MVS subsystem that DB2 uses to control communication and database locking.

IRLM. Internal resource lock manager.

ISO. International Standards Organization.

isolation level. The degree to which a unit of work is isolated from the updating operations of other units of work. See also *cursor stability*, *read stability*, *repeatable read*, and *uncommitted read*.

ISPF. Interactive System Productivity Facility.

ISPF/PDF. Interactive System Productivity Facility/Program Development Facility.

J

Japanese Industrial Standards Committee (JISC). An organization that issues standards for coding character sets.

JCL. Job control language.

JES. MVS Job Entry Subsystem.

JIS. Japanese Industrial Standard.

job control language (JCL). A control language that is used to identify a job to an operating system and to describe the job's requirements.

Job Entry Subsystem (JES). An IBM licensed program that receives jobs into the system and processes all output data that is produced by the jobs.

K

KB. Kilobyte (1024 bytes).

key. A column or an ordered collection of columns identified in the description of a table, index, or referential constraint.

key-sequenced data set (KSDS). A VSAM file or data set whose records are loaded in key sequence and controlled by an index.

KSDS. Key-sequenced data set.

L

labeled duration. A number that represents a duration of years, months, days, hours, minutes, seconds, or microseconds.

large object (LOB). A sequence of bytes representing bit data, single-byte characters, double-byte characters, or a mixture of single- and double-byte characters. A LOB can be up to 2 GB–1 byte in length. See also *BLOB*, *CLOB*, and *DBCLOB*.

leaf page. A page that contains pairs of keys and RIDs and that points to actual data. Contrast with *nonleaf page*.

linkage editor. A computer program for creating load modules from one or more object modules or load modules by resolving cross references among the modules and, if necessary, adjusting addresses.

link-edit. The action of creating a loadable computer program using a linkage editor.

L-lock. Logical lock.

load module. A program unit that is suitable for loading into main storage for execution. The output of a linkage editor.

LOB. Large object.

LOB locator. A mechanism that allows an application program to manipulate a large object value in the database system. A LOB locator is a fullword integer value that represents a single LOB value. An application program retrieves a LOB locator into a host variable and can then apply SQL operations to the associated LOB value using the locator.

LOB table space. A table space that contains all the data for a particular LOB column in the related base table.

local. A way of referring to any object that the local DB2 subsystem maintains. A *local table*, for example, is a table that is maintained by the local DB2 subsystem. Contrast with *remote*.

local subsystem. The unique RDBMS to which the user or application program is directly connected (in the case of DB2, by one of the DB2 attachment facilities).

location. The unique name of a database server. An application uses the location name to access a DB2 database server.

lock. A means of controlling concurrent events or access to data. DB2 locking is performed by the IRLM.

lock duration. The interval over which a DB2 lock is held.

lock escalation. The promotion of a lock from a row, page, or LOB lock to a table space lock because the number of page locks that are concurrently held on a given resource exceeds a preset limit.

locking. The process by which the integrity of data is ensured. Locking prevents concurrent users from accessing inconsistent data.

lock mode. A representation for the type of access that concurrently running programs can have to a resource that a DB2 lock is holding.

lock object. The resource that is controlled by a DB2 lock.

lock promotion. The process of changing the size or mode of a DB2 lock to a higher level.

lock size. The amount of data controlled by a DB2 lock on table data; the value can be a row, a page, a LOB, a partition, a table, or a table space.

log. A collection of records that describe the events that occur during DB2 execution and that indicate their sequence. The information thus recorded is used for recovery in the event of a failure during DB2 execution.

logical index partition. The set of all keys that reference the same data partition.

logical lock (L-lock). The lock type that transactions use to control intra- and inter-DB2 data concurrency between transactions. Contrast with *physical lock (P-lock)*.

logical unit. An access point through which an application program accesses the SNA network in order to communicate with another application program.

logical unit of work (LUW). The processing that a program performs between synchronization points.

logical unit of work identifier (LUWID). A name that uniquely identifies a thread within a network. This name consists of a fully-qualified LU network name, an LUW instance number, and an LUW sequence number.

log initialization. The first phase of restart processing during which DB2 attempts to locate the current end of the log.

log record sequence number (LRSN). A number that DB2 generates and associates with each log record. DB2 also uses the LRSN for page versioning. The LRSNs that a particular DB2 data sharing group generates form a strictly increasing sequence for each DB2 log and a strictly increasing sequence for each page across the DB2 group.

log truncation. A process by which an explicit starting RBA is established. This RBA is the point at which the next byte of log data is to be written.

LRH. Log record header.

LRSN. Log record sequence number.

LU name. Logical unit name, which is the name by which VTAM refers to a node in a network. Contrast with *location name*.

LUW. Logical unit of work.

LUWID. Logical unit of work identifier.

M

mapping table. A table that the REORG utility uses to map between the RIDs of data records in the original copy and in the shadow copy. This table is created by the user.

MB. Megabyte (1 048 576 bytes).

menu. A displayed list of available functions for selection by the operator. A menu is sometimes called a *menu panel*.

migration. The process of converting a DB2 subsystem with a previous release of DB2 to an updated or current release. In this process, you can acquire the functions of the updated or current release without losing the data you created on the previous release.

mixed data string. A character string that can contain both single-byte and double-byte characters.

MPP. Message processing program (in IMS).

MTO. Master terminal operator.

multisite update. Distributed relational database processing in which data is updated in more than one location within a single unit of work.

MVS. Multiple Virtual Storage.

MVS/ESA™. Multiple Virtual Storage/Enterprise Systems Architecture.

MVS/XA™. Multiple Virtual Storage/Extended Architecture.

N

network identifier (NID). The network ID that is assigned by IMS or CICS, or if the connection type is RRSAF, the OS/390 RRS unit of recovery ID (URID).

NID. Network ID.

nonleaf page. A page that contains keys and page numbers of other pages in the index (either leaf or nonleaf pages). Nonleaf pages never point to actual data.

NUL. In C, a single character that denotes the end of the string.

null. A special value that indicates the absence of information.

NUL-terminated host variable. A varying-length host variable in which the end of the data is indicated by the presence of a NUL terminator.

NUL terminator. In C, the value that indicates the end of a string. For character strings, the NUL terminator is X'00'.

O

OASN (origin application schedule number). In IMS, a 4-byte number that is assigned sequentially to each IMS schedule since the last cold start of IMS. The OASN is used as an identifier for a unit of work. In an 8-byte format, the first 4 bytes contain the schedule number and the last 4 bytes contain the number of IMS sync points (*commit points*) during the current schedule. The OASN is part of the NID for an IMS connection.

OBID. Data object identifier.

OS/390. Operating System/390.

P

package. An object containing a set of SQL statements that have been statically bound and that is available for processing. A package is sometimes also called an *application package*.

package list. An ordered list of package names that may be used to extend an application plan.

package name. The name of an object that is created by a BIND PACKAGE or REBIND PACKAGE command. The object is a bound version of a database request module (DBRM). The name consists of a location name, a collection ID, a package ID, and a version ID.

page. A unit of storage within a table space (4 KB, 8 KB, 16 KB, or 32 KB) or index space (4 KB). In a table space, a page contains one or more rows of a table. In a LOB table space, a LOB value can span more than one page, but no more than one LOB value is stored on a page.

page set. Another way to refer to a table space or index space. Each page set consists of a collection of VSAM data sets.

panel. A predefined display image that defines the locations and characteristics of display fields on a display surface (for example, a *menu panel*).

parallel I/O processing. A form of I/O processing in which DB2 initiates multiple concurrent requests for a

single user query and performs I/O processing concurrently (in *parallel*) on multiple data partitions.

parent row. A row whose primary key value is the foreign key value of a dependent row.

parent table. A table whose primary key is referenced by the foreign key of a dependent table.

parent table space. A table space that contains a parent table. A table space containing a dependent of that table is a dependent table space.

participant. An entity other than the commit coordinator that takes part in the commit process. The term participant is synonymous with *agent* in SNA.

partition. A portion of a page set. Each partition corresponds to a single, independently extendable data set. Partitions can be extended to a maximum size of 1, 2, or 4 GB, depending on the number of partitions in the partitioned page set. All partitions of a given page set have the same maximum size.

partitioned page set. A partitioned table space or an index space. Header pages, space map pages, data pages, and index pages reference data only within the scope of the partition.

partitioned table space. A table space that is subdivided into parts (based on index key range), each of which can be processed independently by utilities.

partner logical unit. An access point in the SNA network that is connected to the local DB2 subsystem by way of a VTAM conversation.

piece. A data set of a nonpartitioned page set.

plan. See *application plan*.

plan allocation. The process of allocating DB2 resources to a plan in preparation for execution.

plan name. The name of an application plan.

point of consistency. A time when all recoverable data that an application accesses is consistent with other data. The term point of consistency is synonymous with *sync point* or *commit point*.

precompilation. A processing of application programs containing SQL statements that takes place before compilation. SQL statements are replaced with statements that are recognized by the host language compiler. Output from this precompilation includes source code that can be submitted to the compiler and the database request module (DBRM) that is input to the bind process.

prefix. A code at the beginning of a message or record.

prepare. The first phase of a two-phase commit process in which all participants are requested to prepare for commit.

primary authorization ID. The authorization ID used to identify the application process to DB2.

primary index. An index that enforces the uniqueness of a primary key.

private connection. A communications connection that is specific to DB2.

private protocol access. A method of accessing distributed data by which you can direct a query to another DB2 system. Contrast with *DRDA access*.

private protocol connection. A DB2 private connection of the application process. See also *private connection*.

privilege. The capability of performing a specific function, sometimes on a specific object. The term includes:

- **explicit privileges**, which have names and are held as the result of SQL GRANT and REVOKE statements. For example, the SELECT privilege.
- **implicit privileges**, which accompany the ownership of an object, such as the privilege to drop a synonym one owns, or the holding of an authority, such as the privilege of SYSADM authority to terminate any utility job.

privilege set. For the installation SYSADM ID, the set of all possible privileges. For any other authorization ID, the set of all privileges that are recorded for that ID in the DB2 catalog.

process. In DB2, the unit to which DB2 allocates resources and locks. Sometimes called an *application process*, a process involves the execution of one or more programs. The execution of an SQL statement is always associated with some process. The means of initiating and terminating a process are dependent on the environment.

program. A single compilable collection of executable statements in a programming language.

protected conversation. A VTAM conversation that supports two-phase commit flows.

Q

QMF™. Query Management Facility.

query. A component of certain SQL statements that specifies a result table.

R

RACF. Resource Access Control Facility, which is a component of the SecureWay Security Server for OS/390.

RBA. Relative byte address.

RCT. Resource control table (in CICS attachment facility).

read stability (RS). An isolation level that is similar to repeatable read but does not completely isolate an application process from all other concurrently executing application processes. Under level RS, an application that issues the same query more than once might read additional rows that were inserted and committed by a concurrently executing application process.

rebind. The creation of a new application plan for an application program that has been bound previously. If, for example, you have added an index for a table that your application accesses, you must rebind the application in order to take advantage of that index.

record. The storage representation of a row or other data.

record identifier (RID). A unique identifier that DB2 uses internally to identify a row of data in a table stored as a record. Compare with *row ID*.

record identifier (RID) pool. An area of main storage above the 16-MB line that is reserved for sorting record identifiers during list prefetch processing.

recovery. The process of rebuilding databases after a system failure.

recovery log. A collection of records that describes the events that occur during DB2 execution and indicates their sequence. The recorded information is used for recovery in the event of a failure during DB2 execution.

recovery pending (RECP). A condition that prevents SQL access to a table space that needs to be recovered.

RECP. Recovery pending.

redo. A state of a unit of recovery that indicates that changes are to be reapplied to the DASD media to ensure data integrity.

referential constraint. The requirement that nonnull values of a designated foreign key are valid only if they equal values of the primary key of a designated table.

referential integrity. The state of a database in which all values of all foreign keys are valid. Maintaining referential integrity requires the enforcement of referential constraints on all operations that change the data in a table upon which the referential constraints are defined.

relationship. A defined connection between the rows of a table or the rows of two tables. A relationship is the internal representation of a referential constraint.

relative byte address (RBA). The offset of a data record or control interval from the beginning of the storage space that is allocated to the data set or file to which it belongs.

remote. Any object that is maintained by a remote DB2 subsystem (that is, by a DB2 subsystem other than the local one). A *remote view*, for example, is a view that is maintained by a remote DB2 subsystem. Contrast with *local*.

remote subsystem. Any RDBMS, except the *local subsystem*, with which the user or application can communicate. The subsystem need not be remote in any physical sense, and might even operate on the same processor under the same MVS system.

REORG pending (REORP). A condition that restricts SQL access and most utility access to an object that must be reorganized.

REORP. REORG pending.

repeatable read (RR). The isolation level that provides maximum protection from other executing application programs. When an application program executes with repeatable read protection, rows referenced by the program cannot be changed by other programs until the program reaches a commit point.

request commit. The vote that is submitted to the prepare phase if the participant has modified data and is prepared to commit or roll back.

resource. The object of a lock or claim, which could be a table space, an index space, a data partition, an index partition, or a logical partition.

resource control table (RCT). A construct of the CICS attachment facility, created by site-provided macro parameters, that defines authorization and access attributes for transactions or transaction groups.

resource limit facility (RLF). A portion of DB2 code that prevents dynamic manipulative SQL statements from exceeding specified time limits. The resource limit facility is sometimes called the governor.

resource limit specification table. A site-defined table that specifies the limits to be enforced by the resource limit facility.

restart pending (RESTP). A restrictive state of a page set or partition that indicates that restart (backout) work

needs to be performed on the object. All access to the page set or partition is denied except for access by the:
- RECOVER POSTPONED command
- Automatic online backout (which DB2 invokes after restart if the system parameter LBACKOUT=AUTO)

RESTP. Restart pending.

result table. The set of rows that are specified by a SELECT statement.

RID. Record identifier.

RID pool. Record identifier pool.

RLF. Resource limit facility.

RMID. Resource manager identifier.

RO. Read-only access.

rollback. The process of restoring data changed by SQL statements to the state at its last commit point. All locks are freed. Contrast with *commit*.

ROWID. Row identifier.

row identifier (ROWID). A value that uniquely identifies a row. This value is stored with the row and never changes.

RS. Read stability.

S

SBCS. Single-byte character set.

scalar function. An SQL operation that produces a single value from another value and is expressed as a function name, followed by a list of arguments that are enclosed in parentheses. Contrast with *column function*.

search condition. A criterion for selecting rows from a table. A search condition consists of one or more predicates.

secondary authorization ID. An authorization ID that has been associated with a primary authorization ID by an authorization exit routine.

segmented table space. A table space that is divided into equal-sized groups of pages called segments. Segments are assigned to tables so that rows of different tables are never stored in the same segment.

sequential data set. A non-DB2 data set whose records are organized on the basis of their successive physical positions, such as on magnetic tape. Several of the DB2 database utilities require sequential data sets.

server. The target of a request from a remote requester. In the DB2 environment, the server function is provided by the distributed data facility, which is used to access DB2 data from remote applications.

session. A link between two nodes in a VTAM network.

share lock. A lock that prevents concurrently executing application processes from changing data, but not from reading data. Contrast with *exclusive lock*.

shift-in character. A special control character (X'0F') that is used in EBCDIC systems to denote that the subsequent bytes represent SBCS characters. See also *shift-out character*.

shift-out character. A special control character (X'0E') that is used in EBCDIC systems to denote that the subsequent bytes, up to the next shift-in control character, represent DBCS characters. See also *shift-in character*.

sign-on. A request that is made on behalf of an individual CICS or IMS application process by an attachment facility to enable DB2 to verify that it is authorized to use DB2 resources.

simple table space. A table space that is neither partitioned nor segmented.

single-byte character set (SBCS). A set of characters in which each character is represented by a single byte. Contrast with *double-byte character set* or *multibyte character set*.

SMF. System management facility.

SMS. Storage Management Subsystem.

SNA. Systems Network Architecture.

sourced function. A function that is implemented by another built-in or user-defined function that is already known to the database manager. This function can be a scalar function or a column (aggregating) function; it returns a single value from a set of values (for example, MAX or AVG). Contrast with *built-in function, external function,* and *SQL function*.

source program. A set of host language statements and SQL statements that is processed by an SQL precompiler.

SPUFI. SQL Processor Using File Input.

SQL. Structured Query Language.

SQL authorization ID (SQL ID). The authorization ID that is used for checking dynamic SQL statements in some situations.

SQLCA. SQL communication area.

SQL communication area (SQLCA). A structure that is used to provide an application program with information about the execution of its SQL statements.

SQLDA. SQL descriptor area.

SQL descriptor area (SQLDA). A structure that describes input variables, output variables, or the columns of a result table.

SQL/DS. Structured Query Language/Data System. This product is now obsolete and has been replaced by DB2 for VSE & VM.

SQL escape character. The symbol that is used to enclose an SQL delimited identifier. This symbol is the double quotation mark ("). See also *escape character*.

SQL function. A user-defined function in which the CREATE FUNCTION statement contains the source code. The source code is a single SQL expression that evaluates to a single value. The SQL user-defined function can return only one parameter.

SQL Processor Using File Input (SPUFI). SQL Processor Using File Input. A facility of the TSO attachment subcomponent that enables the DB2I user to execute SQL statements without embedding them in an application program.

SQL return code. Either SQLCODE or SQLSTATE.

SQL string delimiter. A symbol that is used to enclose an SQL string constant. The SQL string delimiter is the apostrophe ('), except in COBOL applications, where the user assigns the symbol, which is either an apostrophe or a double quotation mark (").

SSI. Subsystem interface (in MVS).

SSM. Subsystem member.

stand-alone. An attribute of a program that means it is capable of executing separately from DB2, without using DB2 services.

static SQL. SQL statements, embedded within a program, that are prepared during the program preparation process (before the program is executed). After being prepared, the SQL statement does not change (although values of host variables that are specified by the statement might change).

storage group. A named set of disks on which DB2 data can be stored.

string. See *character string* or *graphic string*.

Structured Query Language (SQL). A standardized language for defining and manipulating data in a relational database.

subsystem. A distinct instance of a relational database management system (RDBMS).

sync point. See *commit point*.

synonym. In SQL, an alternative name for a table or view. Synonyms can be used only to refer to objects at the subsystem in which the synonym is defined.

system administrator. The person at a computer installation who designs, controls, and manages the use of the computer system.

system agent. A work request that DB2 creates internally such as prefetch processing, deferred writes, and service tasks.

system conversation. The conversation that two DB2 subsystems must establish to process system messages before any distributed processing can begin.

Systems Network Architecture (SNA). The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

T

table. A named data object consisting of a specific number of columns and some number of unordered rows. See also *base table* or *temporary table*.

table check constraint. A user-defined constraint that specifies the values that specific columns of a base table can contain.

table space. A page set that is used to store the records in one or more tables.

table space set. A set of table spaces and partitions that should be recovered together for one of these reasons:

- Each of them contains a table that is a parent or descendent of a table in one of the others.
- The set contains a base table and associated auxiliary tables.

A table space set can contain both types of relationships.

task control block (TCB). A control block that is used to communicate information about tasks within an address space that are connected to DB2. An address space can support many task connections (as many as one per task), but only one address space connection. See also *address space connection*.

TB. Terabyte (1 099 511 627 776 bytes).

TCB. Task control block (in MVS).

temporary table. A table that holds temporary data; for example, temporary tables are useful for holding or sorting intermediate results from queries that contain a large number of rows. The two kinds of temporary table, which are created by different SQL statements, are the created temporary table and the declared temporary table. Contrast with *result table*. See also *created temporary table* and *declared temporary table*.

thread. The DB2 structure that describes an application's connection, traces its progress, processes resource functions, and delimits its accessibility to DB2 resources and services. Most DB2 functions execute under a thread structure. See also *allied thread* and *database access thread*.

three-part name. The full name of a table, view, or alias. It consists of a location name, authorization ID, and an object name, separated by a period.

time. A three-part value that designates a time of day in hours, minutes, and seconds.

time duration. A decimal integer that represents a number of hours, minutes, and seconds.

Time-Sharing Option (TSO). An option in MVS that provides interactive time sharing from remote terminals.

timestamp. A seven-part value that consists of a date and time. The timestamp is expressed in years, months, days, hours, minutes, seconds, and microseconds.

TMP. Terminal Monitor Program.

trace. A DB2 facility that provides the ability to monitor and collect DB2 monitoring, auditing, performance, accounting, statistics, and serviceability (global) data.

TSO. Time-Sharing Option.

TSO attachment facility. A DB2 facility consisting of the DSN command processor and DB2I. Applications that are not written for the CICS or IMS environments can run under the TSO attachment facility.

type 1 indexes. Indexes that were created by a release of DB2 before DB2 Version 4 or that are specified as type 1 indexes in Version 4. Contrast with *type 2 indexes*. As of Version 7, type 1 indexes are no longer supported.

type 2 indexes. Indexes that are created on a release of DB2 after Version 6 or that are specified as type 2 indexes in Version 4 or later.

U

UDF. User-defined function.

UDT. User-defined data type. In DB2 for OS/390 and z/OS, the term *distinct type* is used instead of user-defined data type. See *distinct type*.

uncommitted read (UR). The isolation level that allows an application to read uncommitted data.

undo. A state of a unit of recovery that indicates that the changes the unit of recovery made to recoverable DB2 resources must be backed out.

unique constraint. An SQL rule that no two values in a primary key, or in the key of a unique index, can be the same.

unique index. An index which ensures that no identical key values are stored in a table.

unit of recovery. A recoverable sequence of operations within a single resource manager, such as an instance of DB2. Contrast with *unit of work*.

unit of work. A recoverable sequence of operations within an application process. At any time, an application process is a single unit of work, but the life of an application process can involve many units of work as a result of commit or rollback operations. In a *multisite update* operation, a single unit of work can include several *units of recovery*. Contrast with *unit of recovery*.

UR. Uncommitted read.

URID (unit of recovery ID). The LOGRBA of the first log record for a unit of recovery. The URID also appears in all subsequent log records for that unit of recovery.

user-defined data type (UDT). See *distinct type*.

user-defined function (UDF). A function that is defined to DB2 by using the CREATE FUNCTION statement and that can be referenced thereafter in SQL statements. A user-defined function can be an *external function,* a *sourced function,* or an *SQL function.* Contrast with *built-in function*.

UT. Utility-only access.

V

value. The smallest unit of data that is manipulated in SQL.

variable. A data element that specifies a value that can be changed. A COBOL elementary data item is an example of a variable. Contrast with *constant*.

varying-length string. A character or graphic string whose length varies within set limits. Contrast with *fixed-length string*.

version. A member of a set of similar programs, DBRMs, packages, or LOBs.

A version of a program is the source code that is produced by precompiling the program. The program version is identified by the program name and a timestamp (consistency token).

A version of a DBRM is the DBRM that is produced by precompiling a program. The DBRM version is identified by the same program name and timestamp as a corresponding program version.

A version of a package is the result of binding a DBRM within a particular database system. The

package version is identified by the same program name and consistency token as the DBRM.
A version of a LOB is a copy of a LOB value at a point in time. The version number for a LOB is stored in the auxiliary index entry for the LOB.

view. An alternative representation of data from one or more tables. A view can include all or some of the columns that are contained in tables on which it is defined.

Virtual Storage Access Method (VSAM). An access method for direct or sequential processing of fixed- and varying-length records on direct access devices. The records in a VSAM data set or file can be organized in logical sequence by a key field (key sequence), in the physical sequence in which they are written on the data set or file (entry-sequence), or by relative-record number.

Virtual Telecommunications Access Method (VTAM). An IBM licensed program that controls communication and the flow of data in an SNA network.

VSAM. Virtual storage access method.

VTAM. Virtual Telecommunication Access Method (in MVS).

W

warm start. The normal DB2 restart process, which involves reading and processing log records so that data under the control of DB2 is consistent. Contrast with *cold start*.

Z

z/OS. An operating system for the eServer product line that supports 64-bit real storage.

Bibliography

DB2 Universal Database Server for OS/390 and z/OS Version 7 product libraries:

DB2 for OS/390 and z/OS

- *DB2 Administration Guide, SC26-9931*
- *DB2 Application Programming and SQL Guide, SC26-9933*
- *DB2 Application Programming Guide and Reference for Java, SC26-9932*
- *DB2 Command Reference, SC26-9934*
- *DB2 Data Sharing: Planning and Administration, SC26-9935*
- *DB2 Data Sharing Quick Reference Card, SX26-3846*
- *DB2 Diagnosis Guide and Reference, LY37-3740*
- *DB2 Diagnostic Quick Reference Card, LY37-3741*
- *DB2 Image, Audio, and Video Extenders Administration and Programming, SC26-9947*
- *DB2 Installation Guide, GC26-9936*
- *DB2 Licensed Program Specifications, GC26-9938*
- *DB2 Master Index, SC26-9939*
- *DB2 Messages and Codes, GC26-9940*
- *DB2 ODBC Guide and Reference, SC26-9941*
- *DB2 Reference for Remote DRDA Requesters and Servers, SC26-9942*
- *DB2 Reference Summary, SX26-3847*
- *DB2 Release Planning Guide, SC26-9943*
- *DB2 SQL Reference, SC26-9944*
- *DB2 Text Extender Administration and Programming, SC26-9948*
- *DB2 Utility Guide and Reference, SC26-9945*
- *DB2 What's New? GC26-9946*
- *DB2 XML Extender for OS/390 and z/OS Administration and Programming*, SC27-9949
- *DB2 Program Directory, GI10-8182*

DB2 Administration Tool

- *DB2 Administration Tool for OS/390 and z/OS User's Guide, SC26-9847*

DB2 Buffer Pool Tool

- *DB2 Buffer Pool Tool for OS/390 and z/OS User's Guide and Reference, SC26-9306*

DB2 DataPropagator™

- *DB2 UDB Replication Guide and Reference, SC26-9920*

Net.Data®

The following books are available at this Web site: http://www.ibm.com/software/net.data/library.html

- *Net.Data Library: Administration and Programming Guide for OS/390 and z/OS*
- *Net.Data Library: Language Environment Interface Reference*
- *Net.Data Library: Messages and Codes*
- *Net.Data Library: Reference*

DB2 PM for OS/390

- *DB2 PM for OS/390 Batch User's Guide, SC27-0857*
- *DB2 PM for OS/390 Command Reference, SC27-0855*
- *DB2 PM for OS/390 Data Collector Application Programming Interface Guide, SC27-0861*
- *DB2 PM for OS/390 General Information, GC27-0852*
- *DB2 PM for OS/390 Installation and Customization, SC27-0860*
- *DB2 PM for OS/390 Messages, SC27-0856*
- *DB2 PM for OS/390 Online Monitor User's Guide, SC27-0858*
- *DB2 PM for OS/390 Report Reference Volume 1, SC27-0853*
- *DB2 PM for OS/390 Report Reference Volume 2, SC27-0854*
- *DB2 PM for OS/390 Using the Workstation Online Monitor, SC27-0859*
- *DB2 PM for OS/390 Program Directory, GI10-8223*

Query Management Facility (QMF)

- *Query Management Facility: Developing QMF Applications, SC26-9579*
- *Query Management Facility: Getting Started with QMF on Windows, SC26-9582*
- *Query Management Facility: High Peformance Option User's Guide for OS/390 and z/OS, SC26-9581*
- *Query Management Facility: Installing and Managing QMF on OS/390 and z/OS, GC26-9575*

© Copyright IBM Corp. 1983, 2001

- *Query Management Facility: Installing and Managing QMF on Windows, GC26-9583*
- *Query Management Facility: Introducing QMF, GC26-9576*
- *Query Management Facility: Messages and Codes, GC26-9580*
- *Query Management Facility: Reference, SC26-9577*
- *Query Management Facility: Using QMF, SC26-9578*

Ada/370
- *IBM Ada/370 Language Reference, SC09-1297*
- *IBM Ada/370 Programmer's Guide, SC09-1414*
- *IBM Ada/370 SQL Module Processor for DB2 Database Manager User's Guide, SC09-1450*

APL2®
- *APL2 Programming Guide, SH21-1072*
- *APL2 Programming: Language Reference, SH21-1061*
- *APL2 Programming: Using Structured Query Language (SQL), SH21-1057*

AS/400®

The following books are available at this Web site: www.as400.ibm.com/infocenter
- *DB2 Universal Database for AS/400 Database Programming*
- *DB2 Universal Database for AS/400 Performance and Query Optimization*
- *DB2 Universal Database for AS/400 Distributed Data Management*
- *DB2 Universal Database for AS/400 Distributed Data Programming*
- *DB2 Universal Database for AS/400 SQL Programming Concepts*
- *DB2 Universal Database for AS/400 SQL Programming with Host Languages*
- *DB2 Universal Database for AS/400 SQL Reference*

BASIC
- *IBM BASIC/MVS Language Reference, GC26-4026*
- *IBM BASIC/MVS Programming Guide, SC26-4027*

BookManager® READ/MVS
- *BookManager READ/MVS V1R3: Installation Planning & Customization, SC38-2035*

SAA® AD/Cycle® C/370™
- *IBM SAA AD/Cycle C/370 Programming Guide, SC09-1841*
- *IBM SAA AD/Cycle C/370 Programming Guide for Language Environment/370, SC09-1840*
- *IBM SAA AD/Cycle C/370 User's Guide, SC09-1763*
- *SAA CPI C Reference, SC09-1308*

Character Data Representation Architecture
- *Character Data Representation Architecture Overview, GC09-2207*
- *Character Data Representation Architecture Reference and Registry, SC09-2190*

CICS/ESA
- *CICS/ESA Application Programming Guide, SC33-1169*
- *CICS External Interfaces Guide, SC33-1944*
- *CICS for MVS/ESA Application Programming Reference, SC33-1170*
- *CICS for MVS/ESA CICS-RACF Security Guide, SC33-1185*
- *CICS for MVS/ESA CICS-Supplied Transactions, SC33-1168*
- *CICS for MVS/ESA Customization Guide, SC33-1165*
- *CICS for MVS/ESA Data Areas, LY33-6083*
- *CICS for MVS/ESA Installation Guide, SC33-1163*
- *CICS for MVS/ESA Intercommunication Guide, SC33-1181*
- *CICS for MVS/ESA Messages and Codes, GC33-1177*
- *CICS for MVS/ESA Operations and Utilities Guide, SC33-1167*
- *CICS/ESA Performance Guide, SC33-1183*
- *CICS/ESA Problem Determination Guide, SC33-1176*
- *CICS for MVS/ESA Resource Definition Guide, SC33-1166*
- *CICS for MVS/ESA System Definition Guide, SC33-1164*
- *CICS for MVS/ESA System Programming Reference, GC33-1171*

CICS Transaction Server for OS/390
- *CICS Application Programming Guide, SC33-1687*
- *CICS External Interfaces Guide, SC33-1703*
- *CICS DB2 Guide, SC33-1939*
- *CICS Resource Definition Guide, SC33-1684*

IBM C/C++ for MVS/ESA
- *IBM C/C++ for MVS/ESA Library Reference, SC09-1995*
- *IBM C/C++ for MVS/ESA Programming Guide, SC09-1994*

IBM COBOL
- *IBM COBOL Language Reference, SC26-4769*
- *IBM COBOL for MVS & VM Programming Guide, SC26-4767*

IBM COBOL for OS/390 & VM Programming Guide, SC26-9049

Conversion Guide
- *IMS-DB and DB2 Migration and Coexistence Guide, GH21-1083*

Cooperative Development Environment
- *CoOperative Development Environment/370: Debug Tool, SC09-1623*

DataPropagator NonRelational
- *DataPropagator NonRelational MVS/ESA Administration Guide, SH19-5036*
- *DataPropagator NonRelational MVS/ESA Reference, SH19-5039*

Data Facility Data Set Services
- *Data Facility Data Set Services: User's Guide and Reference, SC26-4388*

Database Design
- *DB2 Design and Development Guide* by Gabrielle Wiorkowski and David Kull, Addison Wesley, ISBN 0-20158-049-7
- *Handbook of Relational Database Design* by C. Fleming and B. Von Halle, Addison Wesley, ISBN 0-20111-434-8

DataHub®
- *IBM DataHub General Information, GC26-4874*

Data Refresher
- *Data Refresher Relational Extract Manager for MVS* GI10-9927

DB2 Connect®
- *DB2 Connect Enterprise Edition for OS/2 and Windows: Quick Beginnings, GC09-2953*
- *DB2 Connect Enterprise Edition for UNIX: Quick Beginnings, GC09-2952*
- *DB2 Connect Personal Edition Quick Beginnings, GC09-2967*
- *DB2 Connect User's Guide, SC09-2954*

DB2 Red Books
- *DB2 UDB Server for OS/390 Version 6 Technical Update, SG24-6108-00*

DB2 Server for VSE & VM
- *DB2 Server for VM: DBS Utility, SC09-2394*
- *DB2 Server for VSE: DBS Utility, SC09-2395*

DB2 Universal Database for UNIX, Windows®, OS/2®
- *DB2 UDB Administration Guide: Planning, SC09-2946*
- *DB2 UDB Administration Guide: Implementation, SC09-2944*
- *DB2 UDB Administration Guide: Performance, SC09-2945*
- *DB2 UDB Administrative API Reference, SC09-2947*
- *DB2 UDB Application Building Guide, SC09-2948*
- *DB2 UDB Application Development Guide, SC09-2949*
- *DB2 UDB CLI Guide and Reference, SC09-2950*
- *DB2 UDB SQL Getting Started, SC09-2973*
- *DB2 UDB SQL Reference Volume 1, SC09-2974*
- *DB2 UDB SQL Reference Volume 2, SC09-2975*

Device Support Facilities
- *Device Support Facilities User's Guide and Reference, GC35-0033*

DFSMS

These books provide information about a variety of components of DFSMS, including DFSMS/MVS®, DFSMSdfp™, DFSMSdss™, DFSMShsm, and MVS/DFP™.
- *DFSMS/MVS: Access Method Services for the Integrated Catalog, SC26-4906*
- *DFSMS/MVS: Access Method Services for VSAM Catalogs, SC26-4905*
- *DFSMS/MVS: Administration Reference for DFSMSdss, SC26-4929*
- *DFSMS/MVS: DFSMShsm Managing Your Own Data, SH21-1077*
- *DFSMS/MVS: Diagnosis Reference for DFSMSdfp, LY27-9606*
- *DFSMS/MVS Storage Management Library: Implementing System-Managed Storage, SC26–3123*
- *DFSMS/MVS: Macro Instructions for Data Sets, SC26-4913*
- *DFSMS/MVS: Managing Catalogs, SC26-4914*
- *DFSMS/MVS: Program Management, SC26-4916*
- *DFSMS/MVS: Storage Administration Reference for DFSMSdfp, SC26-4920*
- *DFSMS/MVS: Using Advanced Services, SC26-4921*

- *DFSMS/MVS: Utilities, SC26-4926*
- *MVS/DFP: Using Data Sets, SC26-4749*

DFSORT
- *DFSORT Application Programming: Guide, SC33-4035*

Distributed Relational Database Architecture™
- *Data Stream and OPA Reference, SC31-6806*
- *IBM SQL Reference, SC26-8416*
- *Open Group Technical Standard*

  The Open Group presently makes the following DRDA books available through its Web site at: www.opengroup.org
  - *DRDA Version 2 Vol. 1: Distributed Relational Database Architecture (DRDA)*
  - *DRDA Version 2 Vol. 2: Formatted Data Object Content Architecture*
  - *DRDA Version 2 Vol. 3: Distributed Data Management Architecture*

Domain Name System
- *DNS and BIND, Third Edition, Paul Albitz and Cricket Liu, O'Reilly, ISBN 1-56592-512-2*

Education
- *IBM Dictionary of Computing, McGraw-Hill, ISBN 0-07031-489-6*
- *1999 IBM All-in-One Education and Training Catalog, GR23-8105*

Enterprise System/9000® and Enterprise System/3090™
- *Enterprise System/9000 and Enterprise System/3090 Processor Resource/System Manager Planning Guide, GA22-7123*

High Level Assembler
- *High Level Assembler for MVS and VM and VSE Language Reference, SC26-4940*
- *High Level Assembler for MVS and VM and VSE Programmer's Guide, SC26-4941*

Parallel Sysplex® Library
- *OS/390 Parallel Sysplex Application Migration, GC28-1863*
- *System/390 MVS Sysplex Hardware and Software Migration, GC28-1862*
- *OS/390 Parallel Sysplex Overview: An Introduction to Data Sharing and Parallelism, GC28-1860*
- *OS/390 Parallel Sysplex Systems Management, GC28-1861*
- *OS/390 Parallel Sysplex Test Report, GC28-1963*
- *System/390 9672/9674 System Overview, GA22-7148*

ICSF/MVS
- *ICSF/MVS General Information, GC23-0093*

IMS
- *IMS Batch Terminal Simulator General Information, GH20-5522*
- *IMS Administration Guide: System, SC26-9420*
- *IMS Administration Guide: Transaction Manager, SC26-9421*
- *IMS Application Programming: Database Manager, SC26-9422*
- *IMS Application Programming: Design Guide, SC26-9423*
- *IMS Application Programming: Transaction Manager, SC26-9425*
- *IMS Command Reference, SC26-9436*
- *IMS Customization Guide, SC26-9427*
- *IMS Install Volume 1: Installation and Verification, GC26-9429*
- *IMS Install Volume 2: System Definition and Tailoring, GC26-9430*
- *IMS Messages and Codes, GC27-1120*
- *IMS Utilities Reference: System, SC26-9441*

ISPF
- *ISPF V4 Dialog Developer's Guide and Reference, SC34-4486*
- *ISPF V4 Messages and Codes, SC34-4450*
- *ISPF V4 Planning and Customizing, SC34-4443*
- *ISPF V4 User's Guide, SC34-4484*

Language Environment®
- *Debug Tool User's Guide and Reference, SC09-2137*

National Language Support
- *IBM National Language Support Reference Manual Volume 2, SE09-8002*

NetView®
- *NetView Installation and Administration Guide, SC31-8043*
- *NetView User's Guide, SC31-8056*

Microsoft ODBC
- *Microsoft ODBC 3.0 Software Development Kit and Programmer's Reference, Microsoft Press, ISBN 1-57231-516-4*

OS/390
- *OS/390 C/C++ Programming Guide, SC09-2362*
- *OS/390 C/C++ Run-Time Library Reference, SC28-1663*

- *OS/390 C/C++ User's Guide, SC09-2361*
- *OS/390 eNetwork Communications Server: IP Configuration, SC31-8513*
- *OS/390 Hardware Configuration Definition Planning, GC28-1750*
- *OS/390 Information Roadmap, GC28-1727*
- *OS/390 Introduction and Release Guide, GC28-1725*
- *OS/390 JES2 Initialization and Tuning Guide, SC28-1791*
- *OS/390 JES3 Initialization and Tuning Guide, SC28-1802*
- *OS/390 Language Environment for OS/390 & VM Concepts Guide, GC28-1945*
- *OS/390 Language Environment for OS/390 & VM Customization, SC28-1941*
- *OS/390 Language Environment for OS/390 & VM Debugging Guide, SC28-1942*
- *OS/390 Language Environment for OS/390 & VM Programming Guide, SC28-1939*
- *OS/390 Language Environment for OS/390 & VM Programming Reference, SC28-1940*
- *OS/390 MVS Diagnosis: Procedures, LY28-1082*
- *OS/390 MVS Diagnosis: Reference, SY28-1084*
- *OS/390 MVS Diagnosis: Tools and Service Aids, LY28-1085*
- *OS/390 MVS Initialization and Tuning Guide, SC28-1751*
- *OS/390 MVS Initialization and Tuning Reference, SC28-1752*
- *OS/390 MVS Installation Exits, SC28-1753*
- *OS/390 MVS JCL Reference, GC28-1757*
- *OS/390 MVS JCL User's Guide, GC28-1758*
- *OS/390 MVS Planning: Global Resource Serialization, GC28-1759*
- *OS/390 MVS Planning: Operations, GC28-1760*
- *OS/390 MVS Planning: Workload Management, GC28-1761*
- *OS/390 MVS Programming: Assembler Services Guide, GC28-1762*
- *OS/390 MVS Programming: Assembler Services Reference, GC28-1910*
- *OS/390 MVS Programming: Authorized Assembler Services Guide, GC28-1763*
- *OS/390 MVS Programming: Authorized Assembler Services Reference, Volumes 1-4, GC28-1764, GC28-1765, GC28-1766, GC28-1767*
- *OS/390 MVS Programming: Callable Services for High-Level Languages, GC28-1768*
- *OS/390 MVS Programming: Extended Addressability Guide, GC28-1769*
- *OS/390 MVS Programming: Sysplex Services Guide, GC28-1771*
- *OS/390 MVS Programming: Sysplex Services Reference, GC28-1772*
- *OS/390 MVS Programming: Workload Management Services, GC28-1773*
- *OS/390 MVS Routing and Descriptor Codes, GC28-1778*
- *OS/390 MVS Setting Up a Sysplex, GC28-1779*
- *OS/390 MVS System Codes, GC28-1780*
- *OS/390 MVS System Commands, GC28-1781*
- *OS/390 MVS System Messages Volume 1, GC28-1784*
- *OS/390 MVS System Messages Volume 2, GC28-1785*
- *OS/390 MVS System Messages Volume 3, GC28-1786*
- *OS/390 MVS System Messages Volume 4, GC28-1787*
- *OS/390 MVS System Messages Volume 5, GC28-1788*
- *OS/390 MVS Using the Subsystem Interface, SC28-1789*
- *OS/390 Security Server External Security Interface (RACROUTE) Macro Reference, GC28-1922*
- *OS/390 Security Server (RACF) Auditor's Guide, SC28-1916*
- *OS/390 Security Server (RACF) Command Language Reference, SC28-1919*
- *OS/390 Security Server (RACF) General User's Guide, SC28-1917*
- *OS/390 Security Server (RACF) Introduction, GC28-1912*
- *OS/390 Security Server (RACF) Macros and Interfaces, SK2T-6700 (OS/390 Collection Kit), SK27-2180 (OS/390 Security Server Information Package)*
- *OS/390 Security Server (RACF) Security Administrator's Guide, SC28-1915*
- *OS/390 Security Server (RACF) System Programmer's Guide, SC28-1913*
- *OS/390 SMP/E Reference, SC28-1806*
- *OS/390 SMP/E User's Guide, SC28-1740*
- *OS/390 Support for Unicode: Using Conversion Services, SC33-7050*
- *OS/390 RMF User's Guide, SC28-1949*
- *OS/390 TSO/E CLISTS, SC28-1973*
- *OS/390 TSO/E Command Reference, SC28-1969*
- *OS/390 TSO/E Customization, SC28-1965*
- *OS/390 TSO/E Messages, GC28-1978*
- *OS/390 TSO/E Programming Guide, SC28-1970*
- *OS/390 TSO/E Programming Services, SC28-1971*
- *OS/390 TSO/E REXX Reference, SC28-1975*
- *OS/390 TSO/E User's Guide, SC28-1968*

- *OS/390 DCE Administration Guide, SC28-1584*
- *OS/390 DCE Introduction, GC28-1581*
- *OS/390 DCE Messages and Codes, SC28-1591*
- *OS/390 UNIX System Services Command Reference, SC28-1892*
- *OS/390 UNIX System Services Messages and Codes, SC28-1908*
- *OS/390 UNIX System Services Planning, SC28-1890*
- *OS/390 UNIX System Services User's Guide, SC28-1891*
- *OS/390 UNIX System Services Programming: Assembler Callable Services Reference, SC28-1899*

IBM Enterprise PL/I for z/OS and OS/390
- *IBM Enterprise PL/I for z/OS and OS/390 Language Reference, SC26-9476*
- *IBM Enterprise PL/I for z/OS and OS/390 Programming Guide, SC26-9473*

OS PL/I
- *OS PL/I Programming Language Reference, SC26-4308*
- *OS PL/I Programming Guide, SC26-4307*

Prolog
- *IBM SAA AD/Cycle Prolog/MVS & VM Programmer's Guide, SH19-6892*

RAMAC and Enterprise Storage Server™
- *IBM RAMAC Virtual Array, SG24-4951*
- *RAMAC Virtual Array: Implementing Peer-to-Peer Remote Copy, SG24-5338*
- *Enterprise Storage Server Introduction and Planning, GC26-7294*

Remote Recovery Data Facility
- *Remote Recovery Data Facility Program Description and Operations, LY37-3710*

Storage Management
- *DFSMS/MVS Storage Management Library: Implementing System-Managed Storage, SC26-3123*
- *MVS/ESA Storage Management Library: Leading a Storage Administration Group, SC26-3126*
- *MVS/ESA Storage Management Library: Managing Data, SC26-3124*
- *MVS/ESA Storage Management Library: Managing Storage Groups, SC26-3125*
- *MVS Storage Management Library: Storage Management Subsystem Migration Planning Guide, SC26-4659*

System/370™ and System/390
- *ESA/370 Principles of Operation, SA22-7200*
- *ESA/390 Principles of Operation, SA22-7201*
- *System/390 MVS Sysplex Hardware and Software Migration, GC28-1210*

System Network Architecture (SNA)
- *SNA Formats, GA27-3136*
- *SNA LU 6.2 Peer Protocols Reference, SC31-6808*
- *SNA Transaction Programmer's Reference Manual for LU Type 6.2, GC30-3084*
- *SNA/Management Services Alert Implementation Guide, GC31-6809*

TCP/IP
- *IBM TCP/IP for MVS: Customization & Administration Guide, SC31-7134*
- *IBM TCP/IP for MVS: Diagnosis Guide, LY43-0105*
- *IBM TCP/IP for MVS: Messages and Codes, SC31-7132*
- *IBM TCP/IP for MVS: Planning and Migration Guide, SC31-7189*

VS COBOL II
- *VS COBOL II Application Programming Guide for MVS and CMS, SC26-4045*
- *VS COBOL II Application Programming: Language Reference, GC26-4047*
- *VS COBOL II Installation and Customization for MVS, SC26-4048*

VS Fortran
- *VS Fortran Version 2: Language and Library Reference, SC26-4221*
- *VS Fortran Version 2: Programming Guide for CMS and MVS, SC26-4222*

VTAM
- *Planning for NetView, NCP, and VTAM, SC31-8063*
- *VTAM for MVS/ESA Diagnosis, LY43-0069*
- *VTAM for MVS/ESA Messages and Codes, SC31-6546*
- *VTAM for MVS/ESA Network Implementation Guide, SC31-6548*
- *VTAM for MVS/ESA Operation, SC31-6549*
- *VTAM for MVS/ESA Programming, SC31-6550*
- *VTAM for MVS/ESA Programming for LU 6.2, SC31-6551*
- *VTAM for MVS/ESA Resource Definition Reference, SC31-6552*

Index

Numerics

32K
  option of DSN1COMP utility 564
  option of DSN1COPY utility 572
  option of DSN1PRNT utility 610

A

ABEND
  option of DIAGNOSE utility 126
abnormal-termination
  option of TEMPLATE statement 467
access method services
  new active log definition 541
access path
  RUNSTATS output 431
ACHKP (auxiliary CHECK-pending status) 55, 64
ACTION
  option of DSN1SDMP utility 619
active log
  adding to BSDS 543
  data set
    I/O error 543
  defining in BSDS 541
  deleting from BSDS 541
  enlarging 542
  status 551
AFTER
  option of DSN1SDMP utility 620
ALL
  option of REBUILD INDEX 261
  option of RUNSTATS utility 439
ALLDUMPS, option of DIAGNOSE utility 125
ANCHOR, option of DSN1CHKR utility 556
archive log
  BSDS 543
  deleting 543
ARCHLOG, option of REPORT utility 423
ASCII
  option of LOAD utility 163
  option of UNLOAD utility 482
authorization ID
  naming convention xvi
  secondary
    privileges 7
  SQL
    privileges exercised by 8
AUXERROR, option of CHECK DATA utility 55
auxiliary CHECK-pending status (ACHKP)
  CHECK DATA utility 55, 64
auxiliary index, reorganizing after loading data 212
auxiliary warning status (AUXW)
  CHECK DATA utility 55, 64
AUXW (auxiliary warning status)
  CHECK DATA utility 55, 64
availability
  recovering
    error range 286

B

BACKOUT, option of DSNJU003 utility 538
BASE, option of LISTDEF statement 143
binding
  RUNSTATS output 431, 447
BLOB
  option of LOAD utility 181
  option of UNLOAD utility 501
BSDS (bootstrap data set)
  determining log inventory contents 552
  updating 531
BUFNO, option of TEMPLATE statement 469

C

CANCEL
  option of DSNJU003 utility 537
CARD column
  SYSTABLEPART catalog table
    use by RUNSTATS 449
  SYSTABSTATS catalog table
    use by RUNSTATS 447
CARDF column
  SYSCOLDIST catalog table
    description 448
  SYSINDEXPART catalog table
    use by RUNSTATS 451
  SYSTABLEPART catalog table
    use by RUNSTATS 451
  SYSTABLES catalog table
    use by RUNSTATS 447
  SYSTABSTATS catalog table
    use by RUNSTATS 447
catalog, DB2
  order of recovering objects 287
  recovery 290
CATALOG, option of DSNJU003 utility 536
catalog, VSAM, STOSPACE utility 460
catalog tables
  index recreation 290
  order of recovering 288
  SYSCOLDIST
    CARDF column 448
    COLGROUPCOLNO column 448
    COLVALUE column 448
    FREQUENCYF column 448
    NUMCOLUMNS column 448
    STATSTIME column 448
  SYSCOLUMNS
    COLCARDF column 447
    HIGH2KEY column 447
    LOW2KEY column 447

© Copyright IBM Corp. 1983, 2001

catalog tables *(continued)*
SYSCOLUMNS *(continued)*
STATSTIME column 447
SYSCOPY
effects of COPY 93
SYSINDEXES
CLUSTERING column 448
CLUSTERRATIOF column 448
data collected by STOSPACE utility 459
FIRSTKEYCARDF column 448
FULLKEYCARDF column 448
NLEAF column 448
NLEVELS column 448
STATSTIME column 448
updating with STOSPACE utility 459
SYSINDEXPART
CARDF column 451
data collected by STOSPACE utility 459
example of query 318
FAROFFPOSF column 452
LEAFDIST column 452
NEAROFFPOSF column 452
SYSINDEXSTATS
CLUSTERRATIOF column 448
FIRSTKEYCARD column 448
FIRSTKEYCARDF column 448
FULLKEYCARD column 449
FULLKEYCARDF column 449
KEYCOUNT column 449
KEYCOUNTF column 449
NLEAF column 449
NLEVELS column 449
SYSLOBSTATS
FREESPACE column 453
ORGRATIO column 453
SYSSTOGROUP
data collected by STOSPACE utility 459
updating with STOSPACE utility 459
SYSTABLEPART
CARD column 449
CARDF column 451
data collected by STOSPACE utility 459
example of query 366
FARINDREF column 449
NEARINDREF column 449
PAGESAVE column 450
PERCACTIVE column 450
PERCDROP column 450
PQTY column 451
SECQTYI column 451
SPACE column 451
SQTY column 451
SYSTABLES
CARDF column 447
NPAGES column 447
PCTROWCOMP column 447
STATSTIME column 447
SYSTABLESPACE
data collected by STOSPACE utility 459
DSSIZE column 448
NACTIVE column 448
catalog tables *(continued)*
SYSTABLESPACE *(continued)*
NACTIVEF column 448
STATSTIME column 448
updating with STOSPACE utility 459
SYSTABSTATS
CARD column 447
CARDF column 447
NPAGES column 447
CATMAINT utility
description 49
syntax diagram 49
CCSID
option of LOAD utility 164
option of UNLOAD utility 482
change log inventory utility
active logs 541
archive logs 543
authorization required 540
data sets 540
DELETE statement 541, 543, 544
description 531
examples 546
invoking 541
NEWCAT statement 544
NEWLOG statement 541, 543, 544
option descriptions 533
renaming log data sets 545
renaming system data sets 545
syntax diagram 531
SYSIN stream parsing 540
CHANGELIMIT, option of COPY utility 88
CHAR
option of LOAD utility 177
option of UNLOAD utility 493
CHECK
option of DSN1COPY utility 572
CHECK DATA utility
claims and drains 65
description 53
examples
use after LOAD RESUME 212
using DELETE option 66
using exception tables 211
output 53
syntax diagram 54
use after LOAD REPLACE 210
CHECK INDEX utility
description 67
example 73
logical partitions 71
output 71
running on logical partition 68
syntax diagram 68
use after LOAD 212
CHECK LOB utility
claims and drains 79
description 75
example 80
output 75
restarting 79

CHECK LOB utility *(continued)*
syntax diagram 76
terminating 79
CHECK-pending status
after LOAD 191
CHECK DATA utility 63
indoubt referential integrity 210
resetting 210
CHECKP (CHECK-pending status) 63
CHECKPAGE
option of COPY utility 87
CHECKPT, option of DSNJU003 utility 534
CHKPTRBA, option of DSNJU003 utility 537
CLOB (character large object)
option of LOAD utility 181
option of UNLOAD utility 501
CLUSTERING column of SYSINDEXES catalog table
use by RUNSTATS 448
CLUSTERRATIOF column
SYSINDEXES catalog table
description 448
SYSINDEXSTATS catalog table
description 448
COLCARDF column
SYSCOLUMNS catalog table
description 447
cold start
example, creating a conditional restart control record 543
specifying for conditional restart 535
COLGROUPCOLNO column
SYSCOLDIST catalog table
description 448
COLUMN
option of RUNSTATS utility 158
COLUMN, option of RUNSTATS utility 435
COLVALUE column
SYSCOLDIST catalog table
description 448
comment
SYSIN records 540
commit point
DSNU command 35
REPAIR utility 401, 403
restarting 47
compatibility
CHECK DATA utility 65
CHECK INDEX utility 72
CHECK LOB utility 79
COPY utility 103
COPYTOCOPY utility 120
DIAGNOSE utility 128
LOAD utility 207
MERGECOPY utility 228
MODIFY RECOVERY utility 236
MODIFY STATISTICS utility 241
REBUILD INDEX utility 271
RECOVER utility 297
REORG INDEX utility 323
REORG TABLESPACE utility 385
REPAIR utility 413
compatibility *(continued)*
REPORT utility 425
RUNSTATS utility 444
STOSPACE utility 460
UNLOAD utility 518
utilities access description 44
Compatibility
QUIESCE utility 256
RECOVER utility 300
REORG TABLESPACE utility 383
compressed data, UNLOAD utility description 517
compression
estimating DASD savings 563
concurrency
utilities access description 44
with real-time statistics 722
CONCURRENT, option of COPY utility 88
conditional REORG TABLESPACE, example 391
conditional restart
control record
change log inventory utility 537
creating 543
DSNJU003 utility 537
reading 554
status printed by print log map utility 547
CONLIST, option of DSNU command 33
connection-name
naming convention xvi
CONSTANT
option of UNLOAD utility 500
CONTINUE
option of RECOVER utility 285
CONTINUEIF, option of LOAD utility 166
continuous operation
recovering an error range 286
CONTROL
option of DSNU command 33
COPY
option of LISTDEF statement 142
COPY-pending status
COPY utility 87
LOAD utility 209
REORG TABLESPACE utility 388
resetting 209
TERM utility 101
COPY utility
compatibility 103
copying a list of objects 95
description 81
effect on real-time statistics 719
examples
CHANGELIMIT 107
control statement in JCL utility 36
full image copy 92, 104
incremental image copy 93, 104
invoking DFSMS concurrent copy 106, 107
invoking DFSMS concurrent copy using filter 106
multiple image copy 93
REPORTONLY 107
unauthorized access prevention 28
full image copy 92

COPY utility *(continued)*
multiple image copy 93
option descriptions 84
output 81
performance recommendations 100
syntax diagram 82
COPY1, option of DSNJU003 utility 535
COPY1VOL, option of DSNJU003 utility 536
COPY2, option of DSNJU003 utility 535
COPY2VOL, option of DSNJU003 utility 536
COPYDDN
option of COPY utility 86
option of COPYTOCOPY utility 114
option of LOAD utility 157
option of MERGECOPY utility 224
option of REORG TABLESPACE utility 337
COPYDDN option
LOAD utility 194
REORG TABLESPACE utility 373
COPYDSN, option of DSNU command 34
COPYDSN2, option of DSNU command 34
COPYTOCOPY utility
compatibility 120
description 109
examples
full image copy 117
incremental image copy 117
full image copy 117
option descriptions 112
output 109
syntax diagram 110
correlation ID
naming convention xvi
COUNT option
option of LOAD utility 159
option of REBUILD INDEX utility 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 346
option of RUNSTATS utility 436
CREATE, option of DSNJU003 utility 537
CREATE TRIGGER
modifying statement terminator in DSNTEP2 691
modifying statement terminator in DSNTIAD 689
CRESTART, option of DSNJU003 utility 534
CSRONLY, option of DSNJU003 utility 538
CURRENT
option of DSNU command 35
option of REPORT utility 422
cursor-name
option of EXEC SQL statement 132
CYL, option of TEMPLATE statement 470

D

data
loading into tables 194
DATA
option of CHECK DATA utility 54
option of LOAD utility 155
option of REPAIR utility 404, 405, 407
option of UNLOAD utility 478
data compression
dictionary
building 192, 372
number of records needed to fill 192
using again 192
LOAD utility
description 192
KEEPDICTIONARY option 157, 192
REORG TABLESPACE utility
KEEPDICTIONARY option 344
REORG utility
description 192
KEEPDICTIONARY option 192
Data Facility Sort (DFSORT) 356
data set
concatenating copies
UNLOAD utility 510
copying table space in separate jobs 95
discard 211
error 211
naming convention xvii
recovering
partition 284
security 28
space parameter, changing 319, 371
used by utilities
CHECK DATA utility 211
disposition 27
VSAM 533
data sets
change log inventory utility 540
CHECK DATA utility 61
CHECK INDEX utility 70
CHECK LOB utility 78
COPY utility 90
COPYTOCOPY utility 115
DIAGNOSE utility 127
LOAD utility 183
MERGECOPY utility 225
MODIFY RECOVERY utility 234
MODIFY STATISTICS utility 240
QUIESCE utility 253
REBUILD INDEX utility 264
RECOVER utility 283
REORG INDEX utility 316
REORG TABLESPACE utility 361
REPAIR utility 410
REPORT utility 423
RUNSTATS utility 442
STOSPACE utility 458
UNLOAD utility 508
data sharing
real-time statistics 722
data type
specifying with LOAD utility 177
database
DSNDB01 (DB2 directory database) 288
DSNDB06 (DB2 catalog database) 290
limits 627
naming convention xvii

DATABASE
option of LISTDEF statement 140
DATABASE
option of REPAIR utility 408
DATACLAS, option of TEMPLATE statement 469
DATAONLY, option of DSN1LOGP utility 589
DATAWKnn
data set of REORG utility 25
purpose 25
DATE EXTERNAL
option of LOAD utility 180
option of UNLOAD utility 499
DB2I (DB2 Interactive)
invoking utilities 28
DB2I, option of DSNU command 33
DBCLOB
option of UNLOAD utility 502
DBCLOB (double-byte character large object)
option of LOAD utility 181
DBD statement of REPAIR utility 407
DBD01 directory table space
incremental image copy not allowed 93
MERGECOPY restrictions 223, 228
order of recovering 287, 288
DBID
option of DSN1LOGP utility 590
option of REPAIR utility 408
DBRM (database request module)
member naming convention xvii
partitioned data set naming convention xvii
DD statements
data sets 24
DB2 utility 36
DDF (distributed data facility)
option of DSNJU003 utility 534
ddname, naming convention xvii
DDNAME, option of DSNJU004 utility 547
DEADLINE
option of REORG INDEX utility 308
option of REORG TABLESPACE utility 339
DECIMAL
option of LOAD utility 178
DECIMAL EXTERNAL
option of LOAD utility 178
option of UNLOAD utility 498
DECIMAL PACKED
option of UNLOAD utility 497
DECIMAL PACKED, option of LOAD utility 178
DECIMAL ZONED
option of LOAD utility 178
option of UNLOAD utility 498
declared temporary table
REPAIR utility 408
utility compatibility 8
DEFAULTIF
option of LOAD utility 182
DELAY
option of REORG INDEX utility 310
option of REORG TABLESPACE utility 341
DELETE
option of CHECK DATA utility 56
DELETE *(continued)*
option of DSNJU003 utility 534
option of MODIFY RECOVERY utility 233
option of MODIFY STATISTICS utility 239
statement of REPAIR utility
used in LOCATE block 401
DELETE statement of REPAIR utility 405
deleting
active log from BSDS 541
log data sets with errors 543
DFSMS (Data Facility Storage Management Subsystem)
concurrent copy
invoking with COPY utility 88
DFSORT (Data Facility Sort)
allocates data sets for REORG TABLESPACE 356
determining values
SORTDEVT in CHECK INDEX utility 69
SORTDEVT in LOAD utility 166
messages 364
DIAGNOSE, option of REPAIR utility 408
DIAGNOSE utility
compatibility 128
description 123
example 128
option descriptions 125
syntax diagram 123
diagnosis tool, DSN1CHKR utility 557
directory
DBD01
incremental image copy not allowed 93
order of recovering 288
order of recovering
importance 288
objects 287
SYSLGRNX table
effects of COPY 93
SYSUTILX table space 93
DISCARD, option of REORG TABLESPACE utility 349
discard data set, specifying DD statement for LOAD utility 165
DISCARDDN
option of LOAD utility 165, 172
option of REORG TABLESPACE utility 349
DISCARDS, option of LOAD utility 165
DISCDSN, option of DSNU command 34
DISP, option of TEMPLATE statement 467
DISPLAY
option of DIAGNOSE utility 125
DISPLAY DATABASE command
displaying range of pages in error 286
DISPLAY UTILITY command
monitoring utility status 43
displaying
status of
DB2 utilities 43
DL/I
loading data 193
DPropNR (DataPropagator NonRelational)
options 193
DRAIN
option of REORG INDEX utility 309

DRAIN *(continued)*
option of REORG TABLESPACE utility 341
DRAIN_WAIT
option of REORG INDEX utility 308
option of REORG TABLESPACE utility 339
DROP
option of REPAIR utility 408
DSN
option of TEMPLATE statement 465
DSN1CHKR utility
authorization 557
data sets required 557
description 555
examples 558
option descriptions 555
restrictions 557
syntax diagram 555
DSN1COMP utility
authorization required 566
compression calculations 568
data sets required 566
description 563
identical data rows 568
interpreting output 568
option descriptions 563
output example 569
recommendations 567
sample JCL 568
savings estimate 567
syntax diagram 563
DSN1COPY utility
authorization required 576
copying identity column tables 584
copying tables to other subsystems 584
data sets required 577
description 571
determine data set size 581
determine page size 581
example 579
JCL sample 585
multiple data set table spaces 583
option descriptions 572
preventing inconsistent data 582
printing data sets 584
recommendation 581
resetting log RBA 583
restoring indexes 583
restoring table spaces 584
restrictions 581
syntax diagram 572
tasks 582
translating internal identifiers 582
using image copy input 583
DSN1LOGP utility
data sharing example 598
description 587
example 596
JCL
requirements 594
output example 599
DSN1PRNT utility
authorization required 614
description 609
determine data set size 615
determine page size 615
JCL sample 615
option descriptions 610
recommendations 615
required data sets 614
syntax diagram 610
DSN1SDMP utility
authorization required 620
description 617
JCL sample 623
option descriptions 617
output 624
required data sets 621
syntax diagram 617
DSNACCAV stored procedure
description 649
option descriptions 650
output 653
sample JCL 652
syntax diagram 650
DSNACCOR stored procedure
description 658
example call 669
option descriptions 660
output 673
syntax diagram 660
DSNACCQC stored procedure
description 642
option descriptions 643
output 647
sample JCL 647
syntax diagram 642
DSNAME
option of DSNJU003 utility 534
DSNDB01 database
RECOVER utility access 288
DSNDB06 database
RECOVER utility access 288
DSNJLOGF utility
data sets required 529
description 529
example 529
output 529
DSNJU003 utility
authorization required 540
examples 546
option descriptions 533
syntax diagram 531
DSNJU004 utility
authorization required 548
description of output 549
example 549
option descriptions 547
recommendations 549
running 549
syntax diagram 547
DSNTEJ1 sample 358

DSNTEP2 sample program
how to run 685
parameters 685
program preparation 685
DSNTIAD sample program
how to run 685
parameters 685
program preparation 685
specifying SQL terminator 689
DSNTIAUL sample program
how to run 685
parameters 685
program preparation 685
DSNU CLIST command, syntax diagram 32
DSNU command of TSO
description 31
editing generated JCL 37
example 37
invoking utilities 31
options 32
output 36
DSNUM
option of COPY utility 85
option of COPYTOCOPY utility 112
option of MERGECOPY utility 223
option of MODIFY RECOVERY utility 232
option of RECOVER utility 278
option of REPORT utility 422
DSNUPROC JCL procedure 38
syntax diagram 38
DSNUTILS stored procedure
authorization required 632
data sets 632
description 631
option descriptions 634
output 641
sample JCL 641
syntax diagram 634
DSSIZE
option of DSN1COMP utility 564
option of DSN1COPY utility 573
option of DSN1PRNT utility 611
DSSIZE column
SYSTABLESPACE catalog table
use by RUNSTATS 448
DSSPRINT
data set of COPY utility 25
purpose 25
DUMP
option of DSN1CHKR utility 555
statement of REPAIR utility
used in LOCATE block 401
DUMP statement of REPAIR utility 406

E

EBCDIC
option of LOAD utility 163
option of UNLOAD utility 482
EDIT, option of DSNU command 34
edit routine
LOAD utility 151
REORG TABLESPACE utility 343
END
option of DIAGNOSE utility 125
ENDLRSN, option of DSNJU003 utility 536
ENDRBA, option of DSNJU003 utility 535
ENDTIME, option of DSNJU003 utility 535
ENFORCE
option of LOAD utility 191
ENFORCE, option of LOAD utility 164
ERRDDN
option of CHECK DATA utility 57
option of LOAD utility 164
error
range recovery 286
ERROR RANGE, option of RECOVER utility 281
ESA data compression, estimating DASD savings 563
estimating size of shadow data sets
REORG INDEX utility 315
REORG TABLESPACE utility 359
EVENT
option of OPTIONS statement 247
exception tables 58, 210
EXCEPTIONS
option of CHECK DATA utility 56
option of CHECK LOB utility 76
EXCLUDE, option of LISTDEF statement 138
EXEC SQL
option descriptions 131
EXEC SQL statement
syntax diagram 131
EXEC statement of DB2 utility 36, 41
executing
DSNU CLIST command 31
utilities
DB2I 28
DSNU CLIST command 31
JCL 38, 41
EXPDL, option of TEMPLATE statement 469

F

fallback
RECOVER utility 298
FARINDREF column of SYSTABLEPART catalog table
use by RUNSTATS 449
FAROFFPOSF column of SYSINDEXPART catalog table
catalog query to retrieve value for 366
use by 452
FASTSWITCH, option of REORG INDEX utility 308, 339
field procedure
LOAD utility 203
FILTER
option of DSN1LOGP utility 593
FILTERDDN, option of COPY utility 87
FIRSTKEYCARD column
SYSINDEXSTATS catalog table 448

FIRSTKEYCARDF column
SYSINDEXES catalog table
description 448
SYSINDEXSTATS catalog table 448
FLOAT
option of LOAD utility 163, 179
option of UNLOAD utility 483, 499
FLOAT EXTERNAL
option of LOAD utility 179
option of UNLOAD utility 499
FOR
option of DSN1SDMP utility 620
FOR EXCEPTION, option of CHECK DATA utility 56
FORCEROLLUP
option of LOAD utility 159
option of REBUILD INDEX utility 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 346
option of RUNSTATS utility 436
format
data by DPropNR 193
FORMAT
option of DSN1CHKR utility 555
option of DSN1PRNT utility 613
option of LOAD utility 162
FORWARD, option of DSNJU003 utility 537
free space
REORG TABLESPACE utility 387
REORG utility 324
FREEPAGE
option of DSN1COMP utility 564
FREESPACE column of SYSLOBSTATS catalog table
description 453
FREQUENCYF column
SYSCOLDIST catalog table
description 448
FREQVAL
option of LOAD utility 159
option of REBUILD INDEX utility 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 346
option of RUNSTATS utility 436
FROM TABLE
option of REORG TABLESPACE utility 349
option of UNLOAD utility 485, 520
FROMCOPY
option of COPYTOCOPY utility 113
option of UNLOAD utility 479, 509
FROMCOPYDDN
option of UNLOAD utility 480
FROMCOPYDDN, option of UNLOAD utility 509
FROMLASTCOPY, option of COPYTOCOPY utility 113
FROMLASTFULLCOPY, option of COPYTOCOPY utility 113
FROMLASTINCRCOPY, option of COPYTOCOPY utility 113
FROMVOLUME
option of COPYTOCOPY utility 113
option of UNLOAD utility 480
FULL, option of COPY utility 87
FULLCOPY
option of DSN1COMP utility 565
option of DSN1COPY utility 573
option of DSN1PRNT utility 613
FULLKEYCARD column
SYSINDEXSTATS catalog table 449
FULLKEYCARDF column
SYSINDEXES catalog table
description 448
SYSINDEXSTATS catalog table 449
function
maximum number in select 628

G

GDGLIMIT
option of TEMPLATE statement 470
GENERIC, option of DSNJU003 utility 539
GRAPHIC
option of LOAD utility 177
option of UNLOAD utility 494
GRAPHIC EXTERNAL
option of LOAD utility 177
option of UNLOAD utility 495

H

HALT
option of OPTIONS statement 247
HASH, option of DSN1CHKR utility 555, 556
HEADER, option of UNLOAD utility 490
hexadecimal-constant
naming convention xvii
hexadecimal-string, naming convention xvii
HIGH2KEY column
SYSCOLUMNS catalog table
description 447
HIGHRBA, option of DSNJU003 utility 534
HISTORY
option of LOAD utility 159
option of REBUILD INDEX utility 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 346
option of RUNSTATS utility 436

I

I/O error
marks active log as TRUNCATED 551
IGNOREFIELDS, option of LOAD utility 170
image copy
COPY utility 81
copying 109
COPYTOCOPY utility 109
full
description 81, 109
incremental
copying 117
description 81, 93
making 93
merging 221

image copy *(continued)*
incremental *(continued)*
taking 93
list of objects 95
making after loading a table 209
making in parallel 81, 109
multiple 93
INCLUDE, option of LISTDEF statement 138
inconsistent data indicator 405
INCRCOPY
option of DSN1COPY utility 573
option of DSN1PRNT utility 613
INCURSOR
option of LOAD utility 155, 173
INDDN
option of LOAD utility 155, 172
index
checking 67, 212
naming convention xvii
organization 317
rebuilding 264
recovering 260
space
recovery 259
storage allocated 459
statistics 317
INDEX
option of LISTDEF statement 141
option of MODIFY STATISTICS utility 239
option of REPORT utility 421
option of RUNSTATS utility 159
INDEX ALL, option of REPORT utility 422
INDEX
option of COPY utility 85
option of COPYTOCOPY utility 112
option of RECOVER utility 278
option of REORG INDEX utility 306
option of REORG TABLESPACE utility 345
option of REPAIR utility 399, 400, 402
option of RUNSTATS utility 435, 438
INDEX NONE, option of REPORT utility 422
indexes
REBUILD INDEX utility 259
RECOVER utility 275
INDEXSPACE
option of COPY utility 85
option of COPYTOCOPY utility 112
option of LISTDEF statement 140
option of MODIFY STATISTICS utility 238
option of RECOVER utility 278
option of REPORT utility 421
INDEXSPACES, option of LISTDEF statement 139
INDEXSPACESTATS
contents 707
real-time statistics table 699
INDEXVAL phase of LOAD utility 209
INDREFLIMIT
option of REORG TABLESPACE utility 342
INDSN
option of DSNU command 33
informational COPY-pending status
COPY utility 87
informational COPY-pending status (ICOPY)
resetting 100
INLCOPY
option of DSN1COPY utility 573
option of DSN1PRNT utility 614
inline COPY
creating with LOAD utility 194
creating with REORG TABLESPACE utility 373
using for UNLOAD utility 509
INTEGER
option of LOAD utility 178
option of UNLOAD utility 497
INTEGER EXTERNAL
option of LOAD utility 178
option of UNLOAD utility 497
integrated catalog facility
COPY utility 85
COPYTOCOPY utility 113
MERGECOPY utility 223
RECOVER TABLESPACE utility 233, 279
REORG TABLESPACE utility 363
STOSPACE utility 459
Interactive System Productivity Facility (ISPF) 28
INTO TABLE
option of LOAD utility 167, 213
ISOLATION
option of UNLOAD utility 484
ISPF (Interactive System Productivity Facility)
utilities panels 28
ITEMERROR
option of OPTIONS statement 247

J

JCL (job control language)
COPY utility 92, 106
COPYTOCOPY utility 117
creating for DB2 utility 28
DSNUPROC utility 31
JES3DD
option of TEMPLATE statement 472
job control language (JCL) 92, 117
JOB statement
DB2 utility 36

K

KEEPDICTIONARY
option of LOAD utility
description 157
description of use 192
option of REORG TABLESPACE utility 344
option of REORG utility
description of use 192
key
foreign
LOAD operation 191
length
maximum 628

key *(continued)*
primary
LOAD operation 191
KEY option
REPAIR utility 402
KEYCARD
option of LOAD utility 159
option of REBUILD INDEX utility 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 346
option of RUNSTATS utility 436
KEYCOUNT column
SYSINDEXSTATS catalog table 449
KEYCOUNTF column
SYSINDEXSTATS catalog table 449

L

LARGE
option of DSN1COMP utility 564
option of DSN1COPY utility 574
option of DSN1PRNT utility 611
large partitioned table spaces
RUNSTATS utility 445
LEAFDIST column of SYSINDEXPART catalog table
description 452
LEAFDISTLIMIT
option of REORG INDEX utility 311
LENGTH
option of REPAIR utility 406
LEVELID
option of REPAIR utility 398
LIB
option of DSNUPROC utility 38
LIMIT
option of UNLOAD utility 491
limits, DB2 627
LIST
option of CHECK INDEX utility 68
option of COPY utility 84
option of COPYTOCOPY utility 112
option of LISTDEF statement 142
option of MERGECOPY utility 222
option of MODIFY RECOVERY utility 232
option of MODIFY STATISTICS utility 238
option of QUIESCE utility 252
option of REBUILD INDEX 261
option of RECOVER utility 277
option of REORG INDEX utility 306
option of REORG TABLESPACE utility 335
option of REPORT utility 421
option of RUNSTATS INDEX utility 438
option of RUNSTATS TABLESPACE utility 434
option of UNLOAD utility 480
LIST option
DSNU command 33
LISTDEF
description 143
option descriptions 138
syntax diagram 136
LISTDEFDD
option of OPTIONS statement 246
LOAD utility
building indexes 200
building indexes in parallel 200
collecting inline statistics 212
compatibility 207
compressing data 192
data conversion 198
description 151
effect on real-time statistics 714
example
one partition 189
replace tables in multi-table table space 187
RESUME YES 213
RESUME YES and ENFORCE NO 213
table replacement 187
improving parallel processing 195
improving performance 196, 197
input to 214
KEEPDICTIONARY option 192
LOAD INTO TABLE options 171
loading DB2 tables 212
LOB column 204
LOG
LOB table space 204
making corrections 208
option descriptions 155
ordering records 186
output 151, 214
performance recommendations 194
restarting 206
RESUME with referential constraints 211
ROWID column 203
syntax diagram 153
varying-length data 186
loading
data
DL/I 193
generated by REORG UNLOAD EXTERNAL 188
generated by UNLOAD 188
large amounts 190
referential constraints 190
unique indexes 186
using a cursor 183
variable-length 151
partitions 189
tables 194
LOB (large object)
option of DSN1COPY utility 574
option of DSN1PRNT utility 611
option of LISTDEF statement 143
LOCALSITE
option of RECOVER utility 280
option of REPORT utility 423
LOCATE INDEX statement of REPAIR utility 403
LOCATE statement of REPAIR utility 401
LOCATE TABLESPACE statement of REPAIR utility 401

LOCATION
option of DSNJU003 utility 538
location name
naming convention xvii
lock
utilities
CHECK DATA utility 65
CHECK INDEX 72
CHECK LOB utility 79
COPY utility 103
COPYTOCOPY utility 120
LOAD utility 207
MERGECOPY utility 228
MODIFY RECOVERY utility 236
MODIFY STATISTICS utility 241
QUIESCE utility 256
REBUILD INDEX utility 271
RECOVER TABLESPACE utility 300
REORG INDEX utility 323
REORG TABLESPACE utility 383, 384, 385
REPAIR utility 413
REPORT utility 425
RUNSTATS utility 444
UNLOAD utility 518
locking
STOSPACE utility 460
utilities access description 44
log
data set
adding 531
deleting 531
printing map 547
printing names 547
record structure
types 592
recovery 537, 538
truncation 546
utilities
DSNJU003 (change log inventory) 531
DSNJU004 (print log map) 547
LOG
option of LOAD utility 161
LOG
option of REORG TABLESPACE utility 336
option of REPAIR utility 398
logical unit name
naming convention xvii
LOGONLY
option of RECOVER utility 279
LONGLOG
option of REORG INDEX utility 310
option of REORG TABLESPACE utility 341
LOW2KEY column
SYSCOLUMNS catalog table
description 447
LRSNEND
option of DSN1LOGP utility 589
LRSNSTART
option of DSN1LOGP utility 589
LUNAME
option of DSNJU003 utility 538
LUWID option
option of DSN1LOGP utility 591

M

MAP
option of DSN1CHKR utility 556
option of REPAIR utility 407
MAPDDN
option of LOAD utility 165
MAPPINGTABLE
option of REORG TABLESPACE utility 340
MAXERR
option of UNLOAD utility 484
MAXPRIME
option of TEMPLATE statement 471
MAXRO
option of REORG INDEX utility 309
option of REORG TABLESPACE utility 340
MAXROWS
option of DSN1COMP utility 565
MB
option of TEMPLATE statement 471
member name
naming convention xvii
MEMBER option
option of DSNJU004 utility 547
MERGECOPY utility
compatibility 228
description 221
example 229
option descriptions 222
output 221
syntax diagram 222
message
CHECK DATA utility 53
CHECK INDEX utility 67
CHECK LOB utility 75
DFSORT utility 364
DSNU command 36
MERGECOPY utility 223
MODIFY RECOVERY utility 231
MODIFY STATISTICS utility 237
QUIESCE utility 251
RECOVER utility 275, 276
REORG INDEX utility 303
REORG TABLESPACE utility 327
REPORT utility 419
RUNSTATS utility 431
STOSPACE utility 457
MESSAGE, option of DIAGNOSE utility 126
message by identifier
DSNI012I 285
DSNJ200I 554
DSNU404I 95
DSNU501I 285
MGMTCLAS
option of TEMPLATE statement 469
MIXED
option of LOAD utility 177, 181

MODELDCB
option of TEMPLATE statement 469
MODIFY RECOVERY utility
compatibility 236
description 231
example 236
syntax diagram 232
MODIFY STATISTICS utility
compatibility 241
description 237
example 241
syntax diagram 238
monitoring
index organization 317
table space organization 317, 365
utility status 43

N

NACTIVE column
SYSTABLESPACE catalog table
use by RUNSTATS 448
NACTIVEF column
SYSTABLESPACE catalog table
use by RUNSTATS 448
naming convention
variables in command syntax xvi
NBRSECND
option of TEMPLATE statement 471
NEARINDREF column
SYSTABLEPART catalog table 449
NEAROFFPOSF column
SYSINDEXPART catalog table 452
catalog query to retrieve value for 366
NEWCAT, option of DSNJU003 utility 534
NEWCOPY
option of MERGECOPY utility 223
NEWLOG
option of DSNJU003 utility 533
statement 541
NGENERIC, option of DSNJU003 utility 539
NLEAF column
SYSINDEXES catalog table
description 448
SYSINDEXSTATS catalog table 449
NLEVELS column
SYSINDEXES catalog table
description 448
SYSINDEXSTATS catalog table 449
NOAUXCHKP
option of REPAIR utility 400
NOAUXWARN
option of REPAIR utility 400
NOCHECKPEND
option of REPAIR utility 400
NOCOPYPEND
option of LOAD utility 161
option of REPAIR utility 400
NODUMPS, option of DIAGNOSE utility 125
non-DB2 utilities
effect on real-time statistics 720
non-select dynamic SQL statement
option of EXEC SQL statement 131
NONE
option of DSNU command 33
NOPAD
option of REORG TABLESPACE utility 347
option of UNLOAD utility 483
NOPASSWD, option of DSNJU003 utility 538
NORCVRPEND
option of REPAIR utility 400
normal-termination
option of TEMPLATE statement 467
NOSUBS
option of LOAD utility 164
option of UNLOAD utility 483
NOSYSREC
option of REORG TABLESPACE utility 336
notices, legal 725
NPAGES column
SYSTABLES catalog table
description 447
SYSTABSTATS catalog table 447
NUCOLUMNS column
SYSCOLDIST catalog table
description 448
NULLIF
option of LOAD utility 181
NUMCOLS
option of LOAD utility 159
option of REBUILD INDEX utility 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 346
option of RUNSTATS utility 436
NUMPARTS
option of DSN1COMP utility 564
option of DSN1COPY utility 574
option of DSN1PRNT utility 611

O

OBID
option of DSN1LOGP utility 590
OBIDXLAT
option of DSN1COPY utility 576
OBJECT
option of REPAIR utility 398
object lists, using pattern-matching expressions 144
object status
advisory
resetting 677
restrictive
resetting 677
OFF
option of OPTIONS statement 247
off-loading
error during 551
OFFLRBA, option of DSNJU003 utility 539
OFFPOSLIMIT
option of REORG TABLESPACE utility 342
OFFSET
option of DSN1LOGP utility 593

OFFSET *(continued)*
option of REPAIR utility 404, 405, 406
OPTIONS
option descriptions 246
OPTIONS statement
syntax diagram 245
ORGRATIO column of SYSLOBSTATS catalog table 453
OUTDDN
option of REPAIR utility 409

P

page
making incremental copies 93
recovering 285
PAGE option
DSN1CHKR utility 556
DSN1LOGP utility 591
RECOVER utility 279, 285
REPAIR utility 402, 403
page set REBUILD-pending status (PSRBD)
description 269
resetting 270
PAGES
option of REPAIR utility 407
PAGESAVE column of SYSTABLEPART catalog table
use by RUNSTATS 450
PAGESIZE
option of DSN1COMP utility 564
option of DSN1COPY utility 573
option of DSN1PRNT utility 610
panel
DB2 UTILITIES 28
PARALLEL
option of COPY utility 87
option of RECOVER utility 280
parameter
utility control statement 24, 528
parsing rules
utility control statements 23, 527
PART
option of CHECK DATA utility 55
option of CHECK INDEX utility 68
option of LOAD utility 171, 189
option of QUIESCE utility 252
option of REBUILD INDEX utility 261
option of REORG INDEX utility 307
option of REORG TABLESPACE utility 335
option of REPAIR utility 399, 400, 402, 403
option of RUNSTATS utility 434, 439
option of UNLOAD utility 479, 509
partitioned table space
define no 190
loading 189
replacing a partition 189
unloading 509
PARTLEVEL, option of LISTDEF statement 142
PASSWORD, option of DSNJU003 utility 538
pattern-matching characters, LISTDEF specification 144
pattern-matching expressions 144
PCTFREE
option of DSN1COMP utility 565
PCTPRIME
option of TEMPLATE statement 471
PCTROWCOMP column
SYSTABLES catalog table
use by RUNSTATS 447
pending status
resetting
advisory 677
restrictive 677
PERCACTIVE column of SYSTABLEPART catalog table
use by RUNSTATS 450
PERCDROP column of SYSTABLEPART catalog table
use by RUNSTATS 450
performance
affected by
I/O activity 366
table space organization 366
COPY utility 100
LOAD utility
improving 194
monitoring
database 459
with the STOSPACE utility 459
RECOVER utility 295
REORG TABLESPACE utility
improving 373
REORG utility
improving 320
RUNSTATS utility 443
PHASE
option of DSNU command 35
phases of execution
initialization 35
termination 35
utilities
CATMAINT 49
CHECK DATA 53
CHECK INDEX 67
CHECK LOB 75
COPY 82
COPYTOCOPY utility 110
description 44
EXEC SQL 131
LISTDEF 136
LOAD 151
MERGECOPY 221
MODIFY RECOVERY 231
MODIFY STATISTICS 237
OPTIONS 245
QUIESCE 251
REBUILD INDEX 259
RECOVER 276
REORG INDEX 303
REORG TABLESPACE 328, 379
REPAIR 397
REPORT 419
RUNSTATS 432
STOSPACE 457

phases of execution *(continued)*
utilities *(continued)*
TEMPLATE 463
UNLOAD 477
PIECESIZ
option of DSN1COPY utility 575
option of DSN1PRNT utility 612
point-in-time recovery
performing 291
populating
tables 194
PORT
option of DSNJU003 utility 539
POSITION
option of UNLOAD utility 491
PQTY column
SYSINDEXPART catalog table
description 453
SYSTABLEPART catalog table
use by RUNSTATS 451
PREFORMAT
option of LOAD utility 155, 171, 197
option of REORG INDEX utility 313
option of REORG TABLESPACE utility 356
preformatting active logs
data sets required 529
description 529
example 529
output 529
PREVIEW
option of DSNU command 35
option of OPTIONS statement 246
PRINT
option of DSN1COPY utility 574
option of DSN1PRNT utility 611
print log map utility
authorization required 548
description 547
description of output 549
example 549
JCL requirements 548
option descriptions 547
output sample 552
recommendations 549
running 549
syntax diagram 547
SYSIN stream parsing 548
Print Log Map utility
BSDS timestamps 550
privilege set of a process 7
problem determination
media damage, avoiding 299
process
privilege set of 7
PROMPT
option of DSNU command 35
PSRBD (page set REBUILD-pending) status
description 269
resetting 270
PUNCHDDN
option of REORG TABLESPACE utility 348
PUNCHDDN *(continued)*
option of UNLOAD utility 481
PUNCHDSN
option of DSNU command 34

Q

qualifier-name naming convention xvii
QUIESCE utility
compatibility 256
description 251
example 258
syntax diagram 251
TABLESPACE option 252

R

RBA (relative byte address)
range printed by print log map 549
range specified in active log 541
RBAEND
option of DSN1LOGP utility 589
RBASTART
option of DSN1LOGP utility 589
RBDP* (REBUILD-pending) status
description 269
RBDP* (REBUILD-pending star) status
resetting 270
RBDP (REBUILD-pending) status
description 269, 297
resetting 209, 270, 297
RC0
option of OPTIONS statement 248
RC4
option of OPTIONS statement 248
RC8
option of OPTIONS statement 248
RCPYDSN1
option of DSNU command 34
RCPYDSN2
option of DSNU command 34
reading
conditional restart control records 554
real-time statistics
accuracy 722
for read-only objects 721
for TEMP table spaces 721
for work file table spaces 721
improving concurrency 722
in data sharing 722
when DB2 externalizes 713
real-time statistics tables
altering 699
contents 701
creating 699
description 699
effect of dropping objects 721
effect of mass delete operations 721
effect of SQL operations 721
INDEXSPACESTATS 699
recovering 722

real-time statistics tables *(continued)*
setting up 699
setting update interval 700
starting 701
TABLESPACESTATS 699
REAL TIME STATS
field of panel DSNTIPO 700
rebinding
recommended after LOAD 209
REBUILD
option of REPAIR utility 409
REBUILD INDEX utility
building indexes in parallel 266
compatibility 271
description 259
effect on real-time statistics 718
example 272
option descriptions 260
performance recommendations 266
re-create index 259
syntax diagram 259
REBUILD-pending status
resetting 270
REBUILD-pending status (RBDP) 269
resetting 209
RECDSN
option of DSNU command 34
record count
REORG TABLESPACE utility 379
RECOVER INDEX utility 290
RECOVER TABLESPACE utility
merges multiple image copies 99
RECOVER utility
compatibility 300
description 275
examples
error range 286
JCL and control statements 301
multiple table spaces 283
single partition 284
single table space 283
hierarchy of dependencies 288
input data sets 283
merges multiple image copies 285
option descriptions 277
output 275
performance recommendations 295
syntax diagram 276
table spaces accessed 288
TABLESPACE
containing LOB data 290
recovery
catalog and directory 287
catalog objects 287
data set
partition 284
database
LOB table space 100
REBUILD INDEX utility 259
RECOVER utility 275
REORG makes image copies invalid 92, 100
recovery *(continued)*
directory objects 287
error range 286
page 285
partial 291
real-time statistics tables 722
reporting information 424
table space
description 283
multiple spaces 283
point in time 264
recovery log
backward log 538
forward log 537
RECOVERY option of REPORT utility 420
RECOVERYDDN
option of COPY utility 86
option of COPYTOCOPY utility 114
option of LOAD utility 158, 194
option of MERGECOPY utility 224
option of REORG TABLESPACE utility 337, 373
RECOVERYSITE
option of RECOVER utility 280
option of REPORT utility 423
RECP (RECOVER-pending) status
description 297
resetting 297
REFERENCE
option of UNLOAD utility 484
referential constraint
loading data 190
REORG
option of DSN1COMP utility 565
REORG INDEX utility
description 303
example 325
option descriptions 306
output 324
syntax diagram 304
REORG TABLESPACE utility
building indexes in parallel 375
compatibility 383, 385
description 327
example 388
LOB table space
avoiding COPY-pending status 380
option descriptions 335
output 387
performance recommendations 373
rebalancing partitions 372
shadow data sets
defining 358
syntax diagram 329
REORG utility
compatibility
REORG INDEX 323
compressing data 192
effect on real-time statistics 716
KEEPDICTIONARY option 192
performance 317
performance recommendations 320

REORG utility *(continued)*
shadow data sets
defining 314
reorganizing
indexes 317
partitions 372
table spaces 317, 365
REPAIR utility
compatibility 413, 414
DBD statement 8, 408
description 407
option descriptions 408
syntax diagram 408
declared temporary table compatibility 8, 408
DELETE statement
description 405
syntax diagram 406
description 397
DUMP statement
description 406
option descriptions 406
syntax diagram 406
example 416
LOCATE INDEX statement 403
LOCATE statement
description 401
syntax diagram 401
LOCATE TABLESPACE statement 401
option descriptions 398
output 397, 415
REPLACE statement
description 404
option descriptions 405
syntax diagram 404
SET INDEX statement
description 399
option descriptions 400
syntax diagram 399
SET TABLESPACE statement
description 399
option descriptions 400
syntax diagram 399
syntax diagram 398
VERIFY statement
description 403
option descriptions 404
syntax diagram 404
REPLACE
option of LOAD utility 156
statement of REPAIR utility
description 404
used in LOCATE block 401
replacing
data in a partition 189
table 187
REPORT
option of LOAD utility 159
option of REBUILD INDEX utility 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 346
option of RUNSTATS utility 436
REPORT utility
compatibility 425
description 419
example 429
option descriptions 420
output 425
syntax diagram 420
SYSIBM.SYSLGRNX directory table 419
table space recovery 424
REPORTONLY
option of COPY utility 89
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 342
RESET
option of DSN1COPY utility 576
option of REPAIR utility 405
resetting
pending status
advisory 677
auxiliary CHECK-pending (ACHKP) 677
CHECK-pending (CHKP) 678
COPY-pending 679
group buffer pool RECOVER-pending (GRECP) 679
informational COPY-pending (ICOPY) 100, 680
page set REBUILD- pending (PSRBD) 270
page set REBUILD pending (PSRBD) 680
REBUILD-pending (RBDP) 270, 297, 680
REBUILD-pending star (RBDP*) 270
RECOVER-pending (RECP) 297, 681
REORG-pending (REORP) 682
restart pending 683
restrictive 677
refresh status
REFRESH-pending (REFP) 682
warning status
auxiliary warning (AUXW) 677
RESPORT, option of DSNJU003 utility 539
restart
cannot restart CHECK DATA 64
cannot restart CHECK INDEX 72
cannot restart CHECK LOB 79
cannot restart MODIFY RECOVERY 235
cannot restart MODIFY STATISTICS 241
cannot restart REPAIR 413
cannot restart REPORT 425
cannot restart RUNSTATS 443
cannot restart STOSPACE 460
conditional
control record governs 554
RESTART
option of DSNU command 34
restarting
utilities
CATMAINT 50
CHECK LOB 79
COPY 102
COPYTOCOPY 119
creating your own JCL 47
data set name and volume serial 48
EXEC statement 41

restarting *(continued)*
utilities *(continued)*
LOAD 206
MERGECOPY 228
methods of restart 46
out of space condition 47
QUIESCE 256
REBUILD INDEX 270
RECOVER 299
REORG INDEX 321
REORG TABLESPACE 380
REPORT 425
RUNSTATS 443
STATISTICS keyword 48
UNLOAD 518
using DB2I 46
using the DSNU CLIST command 47
UTPROC 38
RESUME option
LOAD utility 156, 171
RETPD
option of TEMPLATE statement 469
RETRY
option of REORG INDEX utility 309
option of REORG TABLESPACE utility 340
return code
CHANGELIMIT 99
REUSE
option of LOAD utility 161, 171
option of REBUILD INDEX 261
option of RECOVER utility 280
option of REORG INDEX utility 307
option of REORG TABLESPACE utility 335
RID
option of DSN1CHKR utility 556
option of DSN1LOGP utility 591
option of REPAIR utility 402
ROWID
option of LOAD utility 180
option of REPAIR utility 402
option of UNLOAD utility 501
ROWLIMIT
option of DSN1COMP utility 565
running online utilities
data sharing environment 45
JCL 41
RUNSTATS INDEX utility
syntax diagram 438
RUNSTATS TABLESPACE utility
syntax diagram 432
RUNSTATS utility
compatibility 444
description 431
effect on real-time statistics 719
example 453
large partitioned table spaces 445
option descriptions 434
output 445
performance recommendations 443
recommended after LOAD 209
updating catalog columns 449

S

SAMPLE
option of UNLOAD utility 491
scanning rules
utility control statements 23, 527
SCOPE option
CHECK DATA utility 55, 212
SECQTYI column
SYSINDEXPART catalog table
description 453
SYSTABLEPART catalog table
use by RUNSTATS 451
SEGMENT
option of DSN1COPY utility 573
segmented table space
loading and replacing 187
SELECT
option of DSN1SDMP utility 618
select-statement
option of EXEC SQL statement 132
SELECT statement
example
SYSIBM.SYSTABLESPACE 460
list
maximum number of elements 628
sequential data set
loading data 212
SET INDEX statement of REPAIR utility 399
SET TABLESPACE statement of REPAIR utility 399
shadow data sets
allocating
REORG TABLESPACE PART 360
defining
REORG TABLESPACE utility 358
REORG utility 315
estimating size
REORG INDEX utility 315
REORG TABLESPACE utility 359
shift-in character
LOAD utility 174
shift-out character
LOAD utility 174
SHRLEVEL
option of COPY utility 88
option of LOAD utility 156
option of REORG INDEX utility 307
option of REORG TABLESPACE utility 338
option of RUNSTATS utility 436, 439
option of UNLOAD utility 484
simple table space
loading and replacing 187
SIZE
option of DSNUPROC utility 38
SKIP
option of OPTIONS statement 247
SMALLINT
option of LOAD utility 178
option of UNLOAD utility 496
SORTDATA
option of REORG TABLESPACE utility 336

SORTDEVT
option of CHECK DATA utility 57
option of CHECK INDEX utility 69
option of CHECK LOB utility 77
option of LOAD utility 166
option of REBUILD INDEX 262
option of REORG TABLESPACE utility 356
SORTKEYS
option of LOAD utility 162, 196, 200
option of REBUILD INDEX 262
option of REORG TABLESPACE utility 337
SORTNUM
option of CHECK DATA utility 57
option of CHECK INDEX utility 69
option of CHECK LOB utility 77
option of LOAD utility 166
option of REBUILD INDEX 262
option of REORG TABLESPACE utility 356
SORTOUT data set 25, 184
SORTWKnn data set 25
SPACE
option of TEMPLATE statement 470
space, free 324, 387
SPACE column
SYSINDEXPART catalog table
description 453
SYSTABLEPART catalog table
use by RUNSTATS 451
SQL (Structured Query Language)
limits 627
SQL statement terminator
modifying in DSNTEP2 for CREATE TRIGGER 691
modifying in DSNTIAD for CREATE TRIGGER 689
SQL terminator
specifying in DSNTIAD 689
SQTY column
SYSINDEXPART catalog table
description 453
SYSTABLEPART catalog table
use by RUNSTATS 451
STACK
option of TEMPLATE statement 471
stand-alone utilities
control statements 528
option descriptions 528
START TRACE command
option of DSN1SDMP utility 617
STARTIME, option of DSNJU003 utility 535
STARTRBA
option of DSNJU003 utility 535
state
utility execution 43
statistics
space utilization 317, 365
STATISTICS
option of LOAD utility 158
option of REBUILD INDEX 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 344
statistics history
deleting specific entries 241
statistics history *(continued)*
reasons to delete 240
STATSTIME column
SYSCOLDIST catalog table 448
SYSCOLUMNS catalog table 447
SYSINDEXES catalog table 448
SYSTABLES catalog table 447
SYSTABLESPACE catalog table 448
status
CHECK-pending
resetting 210
COPY-pending, resetting 209
option of TEMPLATE statement 467
page set REBUILD-pending (PSRBD) 269
REBUILD-pending (RBDP) 209, 269
REBUILD-pending star (RBDP*) 269
STOGROUP
option of STOSPACE utility 458
stopping
state of utility execution 43
storage group, DB2
disk space 459
storage allocated 459
STORCLAS
option of TEMPLATE statement 469
stored procedure
DSNACCAV 649
DSNACCOR 658
DSNACCQC 642
DSNUTILS 631
STOSPACE utility
compatibility 460
description 457
example 461
monitoring database performance 459
syntax diagram 457
string
naming convention xvii
STRTLRSN, option of DSNJU003 utility 536
SUBMIT
option of DSNU command 35
subsystem
naming convention xvii
SUBTYPE
option of DSN1LOGP utility 592
SUMMARY
option of DSN1LOGP utility 593
option of REPORT utility 423
SWmmWKnn data set 25
SYMLIST
option of DSNU command 33
syntax diagram
CATMAINT utility 49
change log inventory utility 531
CHECK DATA utility 54
CHECK INDEX utility 68
CHECK LOB utility 76
COPY utility 82
COPYTOCOPY utility 110
DIAGNOSE utility 123
DSN1CHKR utility 555

syntax diagram *(continued)*
DSN1COMP utility 563
DSN1COPY utility 572
DSN1PRNT utility 610
DSN1SDMP utility 617
DSNJU003 utility 531
DSNJU004 utility 547
DSNU CLIST command 32
DSNUPROC JCL procedure 38
DSNUTILS stored procedure 634
EXEC SQL statement 131
LISTDEF statement 136
LOAD utility 153
MERGECOPY utility 222
MODIFY RECOVERY utility 232
MODIFY STATISTICS utility 238
OPTIONS statement 245
print log map utility 547
QUIESCE utility 251
REBUILD INDEX utility 259
RECOVER utility 276
REORG INDEX utility 304
REORG TABLESPACE utility 329
REPAIR utility 398
DBD statement 408
DELETE statement 406
DUMP statement 406
LOCATE statement 401
REPLACE statement 404
SET INDEX statement 399
SET TABLESPACE statement 399
VERIFY statement 404
REPORT utility 420
RUNSTATS INDEX utility 438
RUNSTATS TABLESPACE utility 432
STOSPACE utility 457
TEMPLATE statement 463
UNLOAD utility 477
syntax diagrams, how to read xviii
SYSCOPY
data set of COPY utility 25
option of DSN1LOGP utility 590
SYSDISC data set of LOAD utility 25
estimating size 184
SYSDISC data set of REORG utility 25
SYSERR data set of LOAD utility
estimating size 184
purpose 25
SYSIN data set purpose 25
SYSLGRNX directory table
information from REPORT utility 424
SYSMAP data set of LOAD utility
estimating size 184
purpose 25
SYSPRINT
data set for messages and printed output 25
SYSPUNCH
data set of REORG utility 26
purpose 26
SYSREC data set 26
system
limits 627
SYSTEM
option of DSNU command 35
option of DSNUPROC utility 38
system monitoring
index organization 317
table space organization 317, 365
SYSUT1 data set
description 26
estimating size 184
SYSUTILX directory table space
MERGECOPY restrictions 223, 228
order of recovering 287

T

table
dropping
reclaiming space 369
exception 58, 210
TABLE
option of LISTDEF statement 141
option of RUNSTATS utility 158
TABLE
option of REORG TABLESPACE utility 345
option of RUNSTATS utility 434
table name
naming convention xviii
table space
copying 81
determining when to reorganize 317, 365
loading data into 194
merging copies 221
naming convention xviii
reorganizing
using SORTDATA option of REORG utility 366
utilization 317, 365
segmented
COPY utility 97
LOAD utility 187
statistics 317, 365
storage allocated 459
TABLESPACE
option of CHECK DATA utility 55
option of CHECK INDEX utility 69
option of CHECK LOB utility 76
option of COPY utility 84
option of COPYTOCOPY utility 112
option of LISTDEF statement 140
option of MERGECOPY utility 223
option of MODIFY RECOVERY utility 232
option of MODIFY STATISTICS utility 238
option of QUIESCE utility 252
option of REBUILD INDEX utility 261
option of RECOVER utility 278
option of REORG TABLESPACE utility 335
option of REPAIR utility 399, 400, 401
option of REPORT utility 421
option of RUNSTATS utility 434, 439
option of UNLOAD utility 479

TABLESPACES, option of LISTDEF statement 139
TABLESPACESET
  option of QUIESCE utility 252
  option of REPORT utility 420
TABLESPACESTATS
  contents 701
  real-time statistics table 699
TEMPLATE
  option descriptions 464
TEMPLATE statement
  syntax diagram 463
TEMPLATEDD
  option of OPTIONS statement 247
TERM UTILITY command
  description 45
  effect on
    RECOVER utility 299
    REORG TABLESPACE utility 379
    rerunning LOAD 204
    rerunning UNLOAD 518
    restarting COPY 101
terminating
  state of utility execution 44
  utilities
    CATMAINT 50
    CHECK DATA 64
    CHECK INDEX 72
    CHECK LOB 79
    COPY 101
    description 45
    DIAGNOSE 128
    LOAD 204
    MODIFY RECOVERY 235
    MODIFY STATISTICS 241
    QUIESCE 256
    REBUILD INDEX 270
    RECOVER 299
    REORG INDEX 321
    REORG TABLESPACE 380
    REPAIR 413
    REPORT 425
    RUNSTATS 443
    STOSPACE 460
    UNLOAD 518
TEST option
  REPAIR utility 408
TIME
  option of DSNJU003 utility 539
TIME EXTERNAL
  option of LOAD utility 180
  option of UNLOAD utility 500
TIMEOUT
  option of REORG INDEX utility 310
  option of REORG TABLESPACE utility 341
timestamp
  BSDS 550
TIMESTAMP EXTERNAL
  option of LOAD utility 180
  option of UNLOAD utility 500
TOCOPY option of RECOVER utility 280
TOLASTCOPY option of RECOVER utility 281
TOLASTFULLCOPY option of RECOVER utility 281
TOLOGPOINT option of RECOVER utility 279
TORBA option of RECOVER utility 279
TOSEQNO option of RECOVER utility 281
TOVOLUME option of RECOVER utility 281
TRACEID, option of DIAGNOSE utility 127
trigger
  modifying statement terminator in DSNTEP2 691
  modifying statement terminator in DSNTIAD 689
TRK
  option of TEMPLATE statement 470
TRTCH
  option of TEMPLATE statement 472
TRUNCATE
  option of UNLOAD utility 501
TS1 table space 286
TSO
  CLISTs
    DSNU 31
TYPE
  option of DIAGNOSE utility 125
  option of DSN1LOGP utility 592

U

UID
  DSNUPROC utility 38
  option of DSNU command 35
UNCNT
  option of TEMPLATE statement 470
UNICODE
  option of UNLOAD utility 482
unique index
  loading data 186
UNIT
  option of DSNJU003 utility 536
  option of DSNU command 35
  option of TEMPLATE statement 465
unit of recovery
  in-abort 538
  inflight 538
unit of work
  in-commit 537
  indoubt
    conditional restart 537
UNLDDN
  option of REORG TABLESPACE utility 355
  option of UNLOAD utility 481
UNLOAD
  option of REORG INDEX utility 311
  option of REORG TABLESPACE utility 342
UNLOAD utility
  compatibility 518
  compressed data 517
  converting data types 511
  data type compatibility 511
  description 477
  example
    simple case 520
  generating LOAD statements 517
  option descriptions 478

UNLOAD utility *(continued)*
ordering output columns 509
output 477
output field position 513
output field size 513
restarting 518
selecting source partitions 509
selecting source tables 509
STRIP option 516
syntax diagram 477
TRUNCATE option 516
UNLOAD FROM TABLE options 489
unloading
data
variable-length 477
partitions 509
UPDATE
option of CATMAINT utility 49
option of LOAD utility 159
option of REBUILD INDEX utility 262
option of REORG INDEX utility 311
option of REORG TABLESPACE utility 346
option of RUNSTATS utility 436
URID (unit of recovery ID)
option of DSN1LOGP utility 591
utilities
control statements 24
data set disposition 27
description 7
effect on real-time statistics 714
executing
DB2I 28
DSNU CLIST command 31
JCL 38, 41
phases 44
problems during 44
restart 46
LISTDEF description 143
monitoring and controlling 43
online 23, 41
option descriptions 24
target objects
declared temporary table 8
types
CATMAINT 49
change log inventory (DSNJU003) 531
CHECK DATA 53
CHECK INDEX 67
CHECK LOB 75
COPY 81
COPYTOCOPY 109
DIAGNOSE 123
DSN1CHKR 555
DSN1COMP 563
DSN1COPY 571
DSN1LOGP 587
DSN1PRNT 609
DSN1SDMP 617
LOAD 151
MERGECOPY 221
MODIFY RECOVERY 231
utilities *(continued)*
types *(continued)*
MODIFY STATISTICS 237
preformat active log (DSNJLOGF) 529
print log map (DSNJU004) 547
QUIESCE 251
REBUILD INDEX 259
RECOVER 275
REORG 324
REORG INDEX 303
REORG TABLESPACE 327
REPAIR 397
REPORT 419
RUNSTATS 431
STOSPACE 457
UNLOAD 477
UTILITIES panel 28
UTILITY
option of DSNU command 32
utility-id naming convention xviii
UTPRINmm data set purpose 26
UTPRINT data set purpose 26
UTPROC
option of DSNUPROC utility 38

V

validation routine
LOAD utility 151
REORG TABLESPACE utility 343
VALUE
option of DSN1COPY utility 575
option of DSN1LOGP utility 593
option of DSN1PRNT utility 612
VARCHAR
data type
loading 186
option of LOAD utility 177
option of UNLOAD utility 493
VARGRAPHIC
data type
loading 186
option of LOAD utility 178
option of UNLOAD utility 495
VERIFY
statement of REPAIR utility, used in LOCATE block 401
VERIFY statement of REPAIR utility 403
VERSION
option of REPAIR utility 403
VOLCNT
option of TEMPLATE statement 470
VOLUME
option of DSNU command 36
VOLUMES
option of TEMPLATE statement 470
VSAM (virtual storage access method) 223, 233, 279, 363, 459, 460
catalog 85, 113
data sets 533
VSAMCAT, option of DSNJU003 utility 538

W

WAIT option
  DIAGNOSE utility 126
WARNING
  option of OPTIONS statement 248
WHEN
  option of LOAD utility 173
  option of UNLOAD utility 502
WORKDDN
  option of CHECK DATA utility 57
  option of CHECK INDEX utility 69
  option of CHECK LOB utility 76
  option of LOAD utility 162
  option of MERGECOPY utility 223
  option of REBUILD INDEX 261
  option of REORG INDEX utility 313
  option of REORG TABLESPACE utility 355
WRITE
  option of QUIESCE utility 253

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

DB2 Universal Database for OS/390 and z/OS
Utility Guide and Reference
Version 7

Publication No. SC26-9945-01

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

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

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

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

Please tell us how we can improve this book:

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

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

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
SC26-9945-01

Cut or Fold
Along Line

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

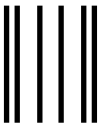

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

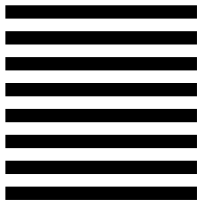

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department HHX/H3
PO BOX 49023
SAN JOSE CA
U. S. A.
95161-9023

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

SC26-9945-01

Cut or Fold
Along Line

Program Number: 5675-DB2

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.
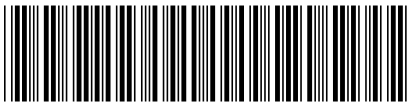
SC26-9945-01